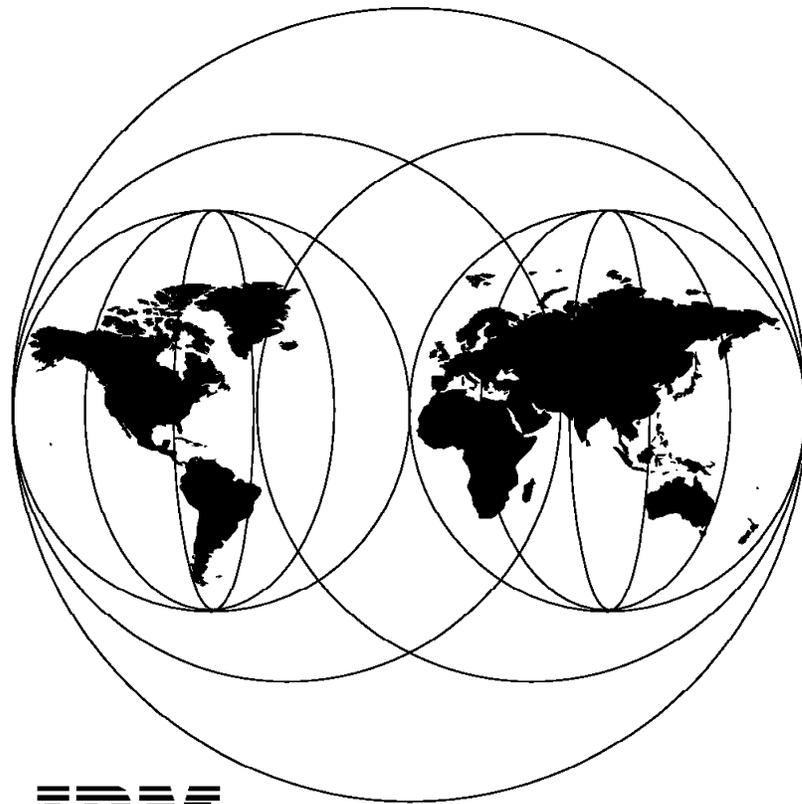# ADSM Concepts

February 1997

IBM

**International Technical Support Organization
San Jose Center**

IBML

International Technical Support Organization

ADSM Concepts

February 1997

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix D, "Special Notices" on page 329.

First Edition (February 1997)

This edition applies to Versions 1 and 2 of ADSM for AIX 5765-203, 5765-564, 5801-AAR, and 5639-393; Versions 1 and 2 of ADSM for MVS 5648-020 and 5655-119; Versions 1 and 2 of ADSM for VM 5648-020 and 5654-A02; Version 2 of ADSM for AS/400 5763-SV2 for OS/400 Version 3 Release 2 and Version 2 of ADSM for AS/400 5716-SV2 for OS/400 Version 3 Release 7; Version 1 of ADSM for HP-UX 5765-304; Version 1 of ADSM for Sun Solaris 5765-303; Version 1 of ADSM for VSE 5686-073; Version 2 of ADSM for Windows NT 5801-AAR, 5807-AAR and 5639-A09

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. QXXE Building 80-E2
650 Harry Road
San Jose, California 95120-6099

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Preface

IBM ADSTAR Distributed Storage Manager (ADSM) is an enterprisewide network storage management solution. This redbook is a tutorial on ADSM and the environments it supports.

ADSM Version 2 was announced in May 1995. Since then enhancements to ADSM have become available through maintenance updates and other products have been enhanced to work with ADSM. This book describes the concepts of ADSM and covers everything that is new.

This redbook is intended for customers, consultants, Business Partners and IBMers who are not familiar with ADSM. It is also useful for people who are familiar with ADSM but want to learn about the latest enhancements.

How This Redbook Is Organized

This redbook contains 352 pages. It is organized as follows:

- Chapter 1, "ADSM Overview"

  This chapter provides an overview of ADSM functions and compares ADSM's backup paradigm with others in the industry.

- Chapter 2, "Disaster Recovery Manager"

  This chapter describes Disaster Recovery Manager and how you can use it to back up your ADSM server.

- Chapter 3, "Tivoli/Plus for ADSM"

  The ADSM Tivoli/Plus module allows you to integrate ADSM operations into a Tivoli environment. This chapter describes TME 10 and the ADSM Tivoli/Plus module and reviews the functions they provide in an ADSM environment.

- Chapter 4, "Andrew File System and Distributed File System Support"

  IBM recently announced ADSM backup support for Andrew File System (AFS) and Distributed File System (DFS) through the AIX 4.1 client. This chapter describes AFS and DFS and the new function ADSM provides to support backup of these file systems.

- Chapter 5, "WebShell Client"

  The ADSM WebShell Client allows you to access an ADSM client through a web browser. This chapter describes the WebShell client and the benefits it provides.

- Chapter 6, "NetTAPE"

  This chapter describes tape management in an AIX environment and shows how NetTAPE and ADSM can simplify operations in that environment.

- Chapter 7, "Network Storage Manager"

  The IBM 3466 Network Storage Manager is an integrated hardware and software storage management solution. ADSM provides the backup function for this product. This chapter describes the Network Storage Manager, its components and the steps required to install and customize it.

- Chapter 8, "Bare Metal Restore"

  Bare metal restore is the process of rebuilding and restoring an ADSM client from the "bare metal" on up. This chapter contains excerpts from the new redbook, *ADSM Client Disaster Recovery: Bare Metal Restore* (SG24-4880).

- Chapter 9, "Miscellaneous Client Enhancements"

  This chapter highlights some of the new functions, platforms, and enhancements to the ADSM clients that have become available through service PTFs.

- Chapter 10, "Miscellaneous Server Enhancements"

  New functions and enhancements to the ADSM server that have shipped through service PTFs are covered in this chapter.

- Chapter 11, "ADSM and Complementary Products"

  This chapter describes the many products that complement ADSM. Topics include products that provide client support (such as Cray UNICOS), products that provide additional function (such as AvailHSM), and products that use ADSM services (such as OnDemand).

- Chapter 12, "ADSM and Databases"

  Databases contain information critical to an enterprise's operation. This chapter contains excerpts from the new edition of the redbook, *Using ADSM to Back Up Databases* (SG24-4335), which describes the major databases and how ADSM can be used to back them up.

- Appendix A, "Server Details"

  This appendix provides details on each ADSM server, including operating system requirements, client platforms, and communication protocols supported.

- Appendix B, "Client Details"

  This appendix provides details on each ADSM client, including operating system requirements and communication protocols supported.

- Appendix C, "Device Support"

  This appendix briefly describes the various ADSM device support modules that are available.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization San Jose Center.

Cybelle Beaulieu is a Distributed Storage Software Specialist at the International Technical Support Organization, San Jose, Center. She writes extensively and teaches IBM classes worldwide on all areas of ADSM. Before joining the ITSO in 1995, Cybelle worked in the Storage Systems Division in San Jose as an ADSM Product Marketing Manager. When she's not working on a redbook, Cybelle is busy chasing after her two toddlers.

Marc Chapel is an ADSM/Tape Libraries Marketing Support Specialist for IBM France and EMEA. He has worked with ADSM for three years in mainframe and UNIX environments and has implemented ADSM for several customers. He holds a PhD in genetics and a degree in opera singing.

Alv Jon Hovda is an AIX Systems Engineer for IBM Norway. He has been with IBM for 27 years and holds an Engineering degree in physics. While in California, Alv fulfilled his dream of driving a Ford Mustang, much to the annoyance of the rest of us, who did not quite fit in the back seat.

Karin Nilsson is a Technical Operations Consultant for IBM Sweden. She has worked for IBM for 17 years with word processing products and MVS. Karin has worked with ADSM on the MVS and AIX platforms since its very beginning, has given several ADSM presentations, and has been involved in may ADSM implementations. Karin tried an aerobics class while in California, and discovered several unknown muscles.

Karl-Ernst Roth is a Systems Engineer for IBM Austria. He has 20 years of experience on IBM's mainframe hardware and software, starting with S/360. He has been working with ADSM for two years and has given many ADSM presentations, and has been involved in many ADSM implementations. After spending six weeks in San Jose, Karl-Ernst was very happy to discover that his six month old son still recognized him and was very excited to see him.

Thanks to the following people for their invaluable contributions to this project:

Maggie Cutler
Technical Editor
International Technical Support Organization, San Jose Center

CD Larson
IBM Storage Systems Division

Steve Luko
IBM Storage Systems Division

Chris Malack
IBM Storage Systems Division

Freddy Saldana
IBM Storage Systems Division

Don Warren
IBM Storage Systems Division

Comments Welcome

We want our redbooks to be as helpful as possible. Should you have any comments about this or other redbooks, please send us a note at the following address:

 redbook@vnet.ibm.com

Your comments are important to us!

Chapter 1. ADSM Overview
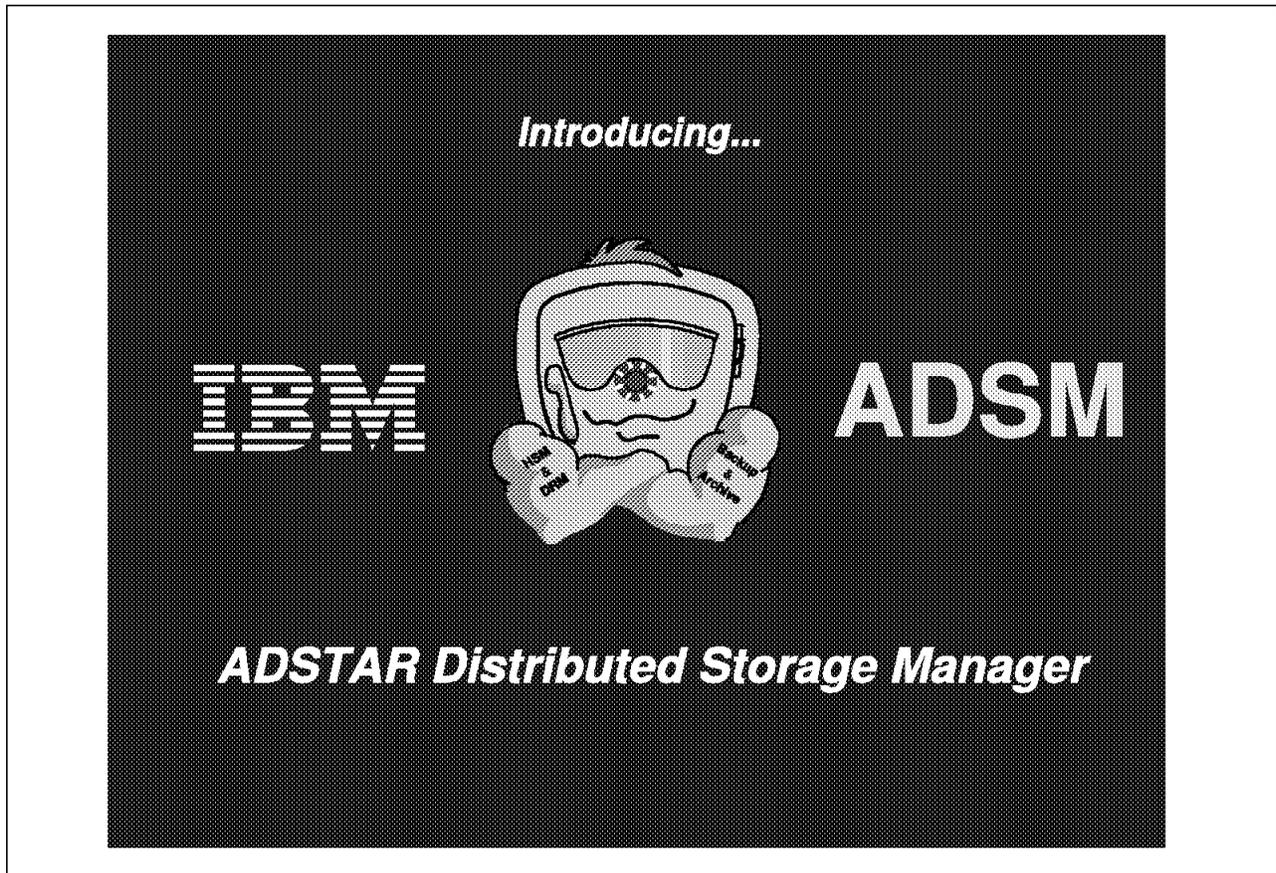
Introducing ... IBM ADSM

ADSM, an IBM client-server solution for distributed data management, supports a wide variety of IBM and non-IBM platforms for both small and large systems. In this book we show you how ADSM is geared toward solving the storage management problems for your entire enterprise. We explain the services of:

- Backup/restore
- Archive/retrieve
- Hierarchical Storage Management (HSM)

and we look at some key functions that make ADSM a robust and flexible solution.

The great flexibility of ADSM's policy management structure enables you to decide how to handle each user's data on the basis of your own defined policies.

You can get your users involved and have them initiate ADSM services, or you can provide ADSM services transparently to your users by centrally automating the ADSM functions.

Trademarks

Owner	Trademark
■ IBM Corporation	■ IBM, AIX, MVS/ESA, VM, OS/2, ADSTAR, VTAM, Lotus Notes, Tivoli, SystemView
■ ADIC	■ ADIC
■ Apple Computer	■ Apple, Macintosh
■ Attachmate Corp.	■ Attachmate, Extral
■ Auspex, Inc.	■ Auspex
■ Bull, Inc.	■ Bull, BOS/X, DPX/2
■ Cray Research, Inc.	■ Cray, UNICOS
■ DataTools	■ DataTools, SQL-BackTrack
■ Digital Equipment Corp.	■ Digital, DEC ULTRIX, Digital UNIX, VMS, OpenVMS
■ DLI	■ DLI
■ ETIUSA, Inc.	■ BackHomel, ETI
■ Exabyte Corp.	■ Exabyte, EXB
■ Fujitsu Ltd.	■ Fujitsu, UXP
■ Hewlett-Packard	■ HP, HP-UX
■ Interlink Computer Sciences Inc.	■ Interlink
■ Microsoft, Inc.	■ Windows NT, Windows, LAN Manager, Named Pipes, Windows 95
■ NCR	■ NCR
■ Nippon Electronics Company	■ NEC
■ Novell, Inc.	■ NetWare, IPX/SPX
■ Odetics, Inc.	■ Odetics, ATL

Trademarks

Trademarks...

Owner	Trademark
■ Oracle, Inc.	■ Oracle
■ The Santa Cruz Operation, Inc.	■ SCO, Open Desktop
■ Sequent	■ Sequent
■ Siemens-Nixdorf	■ Sinix, ADABAS/D
■ Silicon Graphics, Inc.	■ SGI, IRIX, Silicon Graphics
■ Storage Technology, Inc.	■ STK, StorageTek
■ Sun Microsystems	■ Sun, SunOS, Solaris
■ Sybase, Inc.	■ Sybase
■ SAP AG	■ SAP R/3, SAP
■ Tandem Computer	■ Tandem, Guardian
■ Transitional Technologies, Inc.	■ TTI
■ Wang Laboratories, Inc.	■ WangDAT, NetSpace, Avail
■ X/Open, Inc.	■ UNIX

Other trademarks used are the property of their respective owners.

Trademarks...

Storage Management
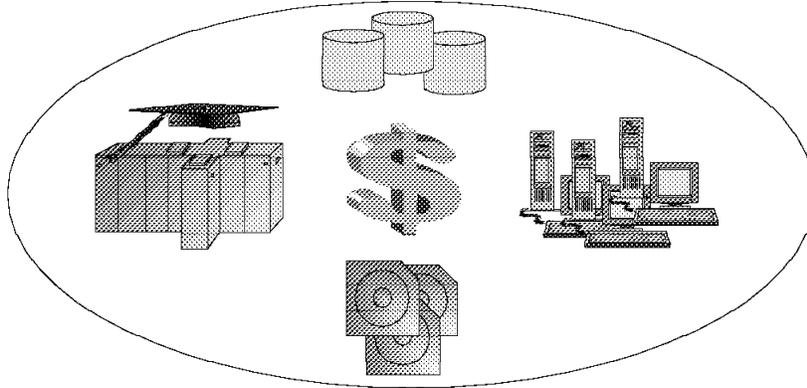
1970s: IBM introduced Hierarchical Storage Manager
1980s: IBM introduced DASD fast write, dual copy
1990s: IBM introduced DFSMS, ABARS, concurrent copy

Storage Management

IBM was first in storage management in 1970 and will continue to provide solutions for storage management beyond the year 2000.

Let us look the evolution of storage management:

- 1970s: IBM introduced Hierarchical Storage Manager
- 1980s: IBM introduced DASD fast write, dual copy
- 1990s: IBM introduced DFSMS, ABARS, concurrent copy

Today's competitive environment requires careful planning and the safeguarding of all company assets, including data on your network. The loss of that data is critical and could mean that you lose your competitiveness, if not your entire business. ADSM saves more than just your data, however. It saves you money, time, and effort. It enables you to invest for the future, and it is a safe buy.

Today's Challenge
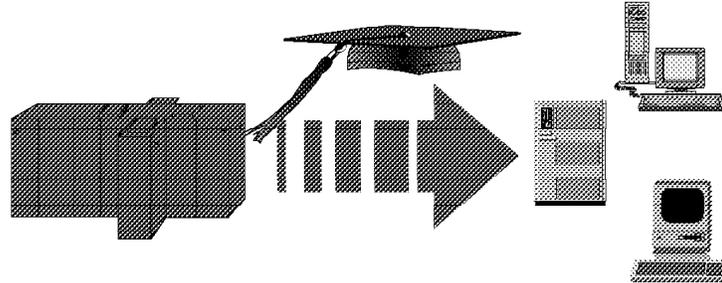
Bring IBM's Industrial-Strength Storage
Management Function to New Computing Worlds.

Today's Challenge

The role of the workstation has expanded tremendously during recent years.
The growth that we are seeing today is similar to the growth experienced in the
centralized mainframe environments of the 1980s. Unfortunately, rapid growth
also brings new levels of complexity to the management and support of the
expanding environments. As the environment grows, it becomes difficult to
coordinate all activities, and we often lose sight of the overall objectives.

Along with complexity we have to consider the cost of the equipment as well as
the hidden expenses involved in the management of the new enterprise.

All of this in the face of the extreme economic pressures of the day!

Recovery Scenarios
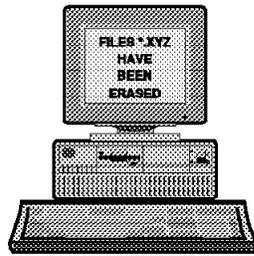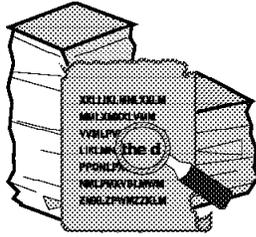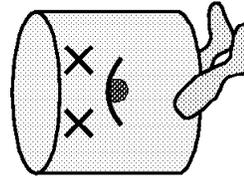
Single File · Hard Disk · Theft · Audit · Disaster

Recovery Scenarios

There are at least five recovery scenarios you should take into account when evaluating what is the best storage management solution for your company.

Causes of Data Loss
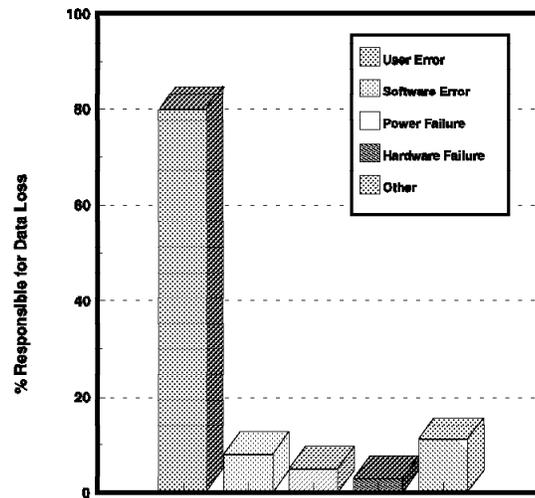
According to a Digital Storage Systems study, the major cause of data loss is operator error; that is, user error. Users may inadvertently save a file that contains erroneous information, or they may erase a critical file by mistake!

Customer Concerns

- ✶ **Multivendor Environments**

- ✶ **Mainframe and Distributed Storage Managed Separately**

- ✶ **Labor-Intensive Tasks and Manual Inventory**

- ✶ **Inconsistent and/or Inadequate Management**

- ✶ **Backup Driven by Exception Conditions**

- ✶ **Storage Management Often Costly and Unreliable**

Customer Concerns

Without doubt, protecting corporate data through automated backup software has become an essential part of managing the network environment.

Multivendor Environments

Today the struggle is with multivendor environments. The struggle is alleviated, however, with ADSM, which supports a broad spectrum of client-server platforms.

Mainframe and Distributed Storage Managed Separately

In the 1970s and 1980s, backup of data within the glass house was well understood and provided through products like IBM's DFSMS/MVS.

In the 1990s, however, more and more critical data resides on workstations throughout an enterprise, and the issues of data management, backup, and availability in a distributed environment must now be addressed.

The explosion of data has created a situation where mainframe storage is managed separately from distributed storage. Furthermore, in a distributed environment, the data is also managed separately, if at all. For example, storage in a UNIX environment may be managed differently from storage in a NetWare environment.

Labor-Intensive Tasks and Manual Inventory

Leaving the protection of critical data to individual users does not necessarily guarantee that data will be consistently backed up or that backup versions will be stored in a safe place. The truth of the matter is that end users often view backing up their data as a nonproductive task that frequently gets put off for another day!

What is more, if users do back up their own data, recovering it can be more labor intensive and time-consuming than they would prefer. Organizations do not have the luxury of losing an hour's worth of data because of hardware failure or human error. When data is lost, the impact is severe and sometimes devastating. Under these circumstances users should not have to spend unnecessary time searching through diskettes and tapes trying to locate the correct version of data they need to recover. With today's backup solutions, users may find the recovery process more difficult and labor intensive than the actual backup process itself!

Inconsistent and/or Inadequate Management

The various backup packages available today could introduce inconsistent data management. Users in one area may back up their data once a month, whereas users in other areas may back up their data once a week. Furthermore, this self-management may not be sufficient to protect the organization's critical data.

Backup Driven by Exception Conditions

Data is only as safe as the last backup. And yet, it is not unusual for users not to back up their data until driven by the painful experience of having to re-create a file lost because of a hard disk failure.

Storage Management Often Costly and Unreliable

Traditionally companies have wasted lots of money by letting users spend a gold mine on their way of buying and managing storage. With ADSM you build on your investment and do not throw money out the window.

ADSM Operating Environments

ADSM Clients

Clients

AIX *	NCR UNIX
AT&T GIS	NEC *
Auspex	NetWare *
Bull BOS/X	OS/400 ***
Bull DPX/2	OS/2 *
Cray UNICOS **	Pyramid Nile
Digital ULTRIX	SCO
Digital UNIX	Sequent
Digital VMS **	SGI
DOS	Sinix *
Fujitsu UXP **	Sun OS *
HP-UX *	Tandem **
Macintosh	Windows *
MS LanManager	Windows NT
MVS OpenEdition *	Windows 95

Communication Protocols

APPC
IPX/SPX
Named Pipes
PWSCS
Shared Memory
TCP/IP
3270 Emulation

ADSM Servers

OS/390	AIX
MVS/ESA	HP-UX
VM/ESA	Solaris
VSE/ESA	Windows NT
OS/400	

Note: Not all clients attach to all servers using all protocols!
* API available
** Through alliance or offering
*** API only

ADSM Operating Environments

As a client-server application, ADSM supports a variety of both IBM and non-IBM platforms and commonly used communication protocols.

ADSM Clients

With ADSM, you can protect data on any of the platforms listed on the left of this chart, the so-called ADSM clients. The functions of the backup-archive client are covered beginning with "Backup" on page 14.

ADSM also provides an application programming interface (API) that enables you to provide online backup services for your own applications (see the clients marked with an asterisk.)

Clients with two asterisks are available through an alliance or offering. For more information refer to "ADSM CRAY UNICOS Client" on page 172, "ADSM Fujitsu UXP/V Client" on page 174, "OpenVMS Client for ADSM" on page 176, and "Tandem Guardian Client" on page 178.

The client with three asterisks (OS/400) uses only the API. There is no backup-archive client currently available.

Communication Protocols

Data can be sent through the communication protocols listed from the ADSM clients to the ADSM server. ADSM communicates at the software level and is not dependent on any network topology. Each ADSM server supports some or all of these communication protocols.

For more information about the specific communication protocols supported for each server, refer to Appendix A, "Server Details" on page 245. For more information about the specific communication protocols supported for each client, refer to Appendix B, "Client Details" on page 283.

ADSM Servers

The arriving data is stored in and managed by one of the available ADSM servers. The functions of the ADSM server are covered in "Server Topics" on page 23.
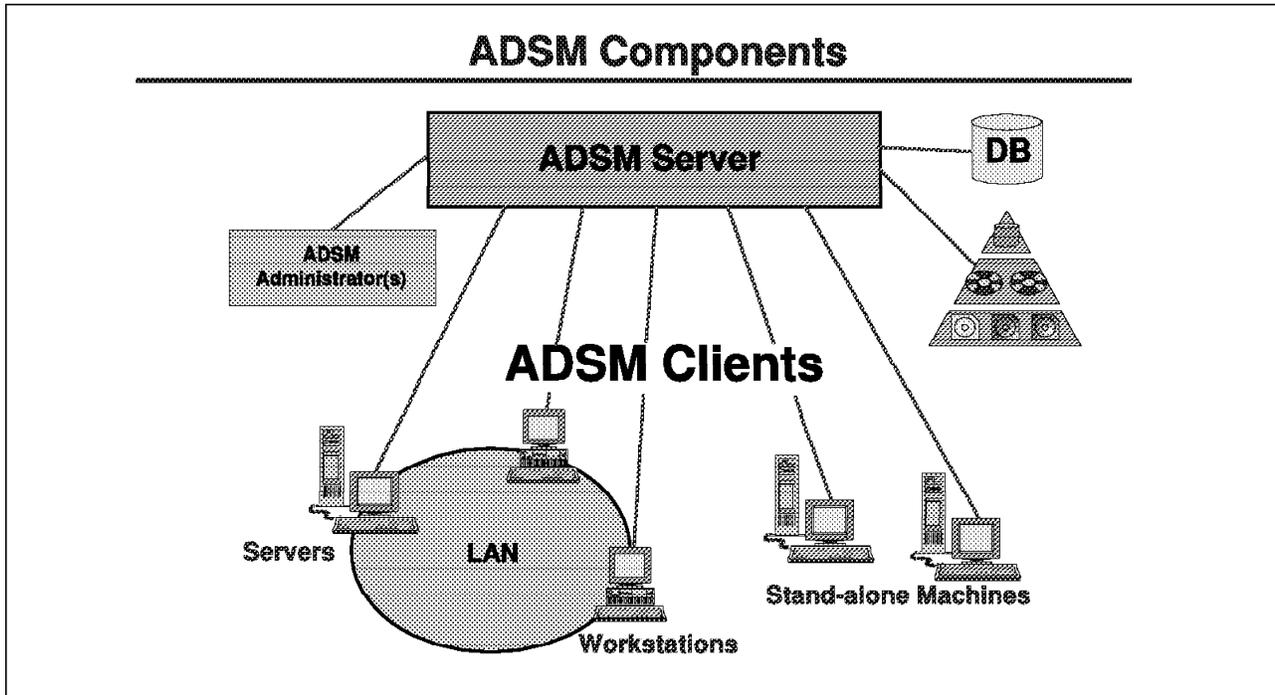
ADSM Components

ADSM Components

ADSM has three major components:

- The ADSM client can be a personal computer, a workstation, or a local area network (LAN) file server. The ADSM client holds the critical data that has to be protected.
- The client's data is sent to the *ADSM server* through a communication protocol. The ADSM server stores the data in ADSM storage pools and uses a database and recovery log to track the location of the data in the storage pools. The server can be a mainframe or a workstation.
- The *ADSM administrator* controls the ADSM server from an administrative client. The administrator also defines the policies that will help protect data in your enterprise. (These policies contain "rules" specifying how data will be treated during a backup or archive operation.)

Note: Throughout this book, when we use the term *client*, we are referring to the ADSM backup-archive client. (This is an important point because the ADSM backup-archive client can be a LAN file server.) When we use the term *server*, we are referring to the ADSM server.
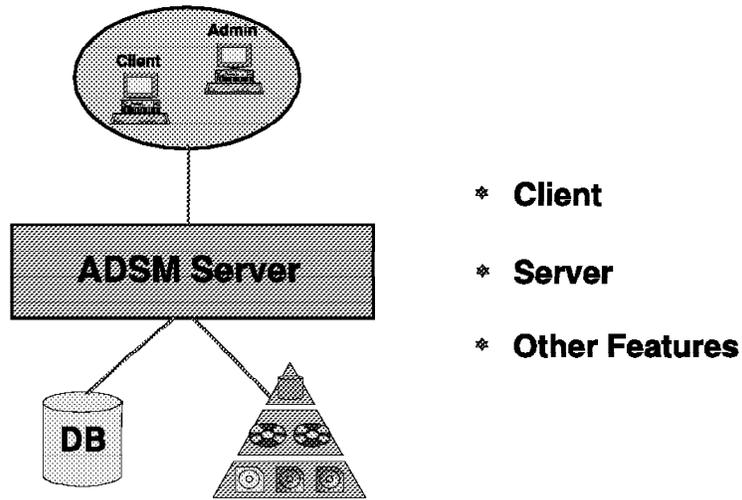
Agenda

Let us now explain in detail how ADSM can work for you. The three major parts of the presentation are:

- Client

- Server

- Other features

Backup

* Incremental or Selective
* Automatically Scheduled
* Multiple Versions and Policy Managed
* User or Help Desk Restore
* Restore by Date and Time Range
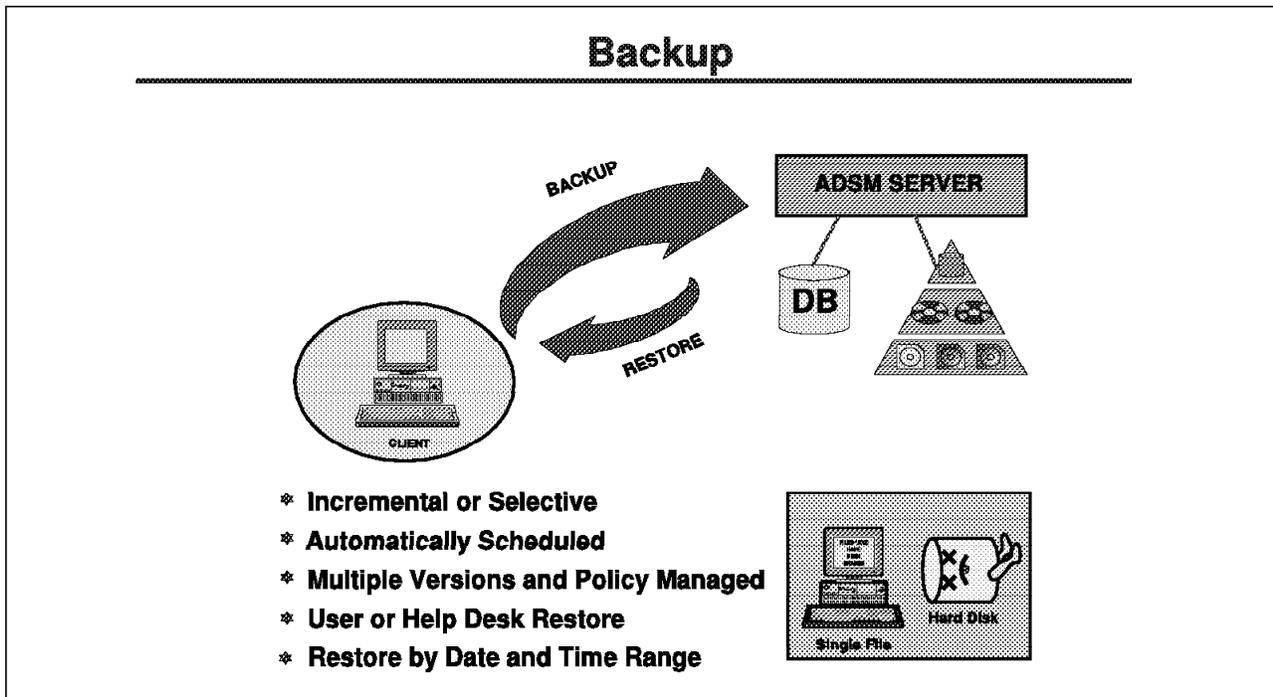
Backup

Backup enables you to save copies of files on the ADSM server. Thus, your data at the ADSM client is protected in the event of data loss due to a hardware or software failure, accidental deletion, and/or logical corruption. You can back up individual files, directories, or subdirectories.

ADSM does not just use the archive bit to tell whether a file has been changed. ADSM checks the modified date and time stamp, the size of the file, and its attributes to see whether it should back up the file.

From a client, you can restore files that have been backed up to the ADSM server. During a restore, ADSM copies the backup version from the ADSM server to your workstation or LAN server.

Restore enables you to recover a single file or an entire file system.

Incremental or Selective

Incremental backup backs up only those files that have changed. This is a true, progressive, incremental backup. In traditional backup products, you must take a periodic full dump to consolidate your backup copies. Because of ADSM′s storage hierarchy management collocation and space reclamation (which we explain in "Space Reclamation" on page 33 and "Storage Pool - Collocation" on page 34) and ADSM′s database, you *never do a full dump.* to the server, and from that point forward, *only changed files are sent.* A thorough discussion of exactly how ADSM gets away with not doing full dumps is described in "Common Backup Paradigm - 1" on page 47.

Incremental backup reduces the amount of data sent on a daily basis and has a major impact on both network utilization and performance because you are sending less data!

Selective backup enables you to explicitly back up few files, whether or not they have changed.

Automatically Scheduled

Backups can be, and usually are, automated. In this way, the user is not burdened with remembering to initiate a backup. The only a user has to do is leave the workstation or server up and running.

The administrator can choose to automate the backup process through the ADSM centralized scheduling feature or by using client command lines, programs, or workstation operating system facilities to initiate the backup process.

Multiple Versions and Policy Managed

ADSM's policy management controls the number of versions of backup files to keep. We cover policy management in detail in "Policy Management" on page 24.

User or Help Desk Restore

If a file is damaged, you can request, without the aid of an administrator, that ADSM restore the current or a specific backup version. You can restore only those files that you have backed up, unless you have been granted authority to another person's saved files.

In response to many customer requirements, the ADSM Version 2 server can now define schedules to perform restores, retrieves, client operating system commands, and macros. Thus an ADSM administrator can act in a help desk role to restore files that users request.

As with all central scheduling, there must be a process or daemon running DSMC SCHED on the client waiting to catch the schedule to be run. This process runs constantly, listening for the server to notify the client that there is a schedule to run. In some cases listening means that the client periodically queries the server for work. In other cases listening means that the server will contact the client when there is a schedule to run.

Restore by Date and Time Range

Sometimes, you may need to recover files on your workstation to a particular date and time, for example, Monday, July 8, at 10:30 a.m. You specify the TODATE and TOTIME options with the desired parameters in the restore command, and ADSM will recover your workstation data to that particular point.
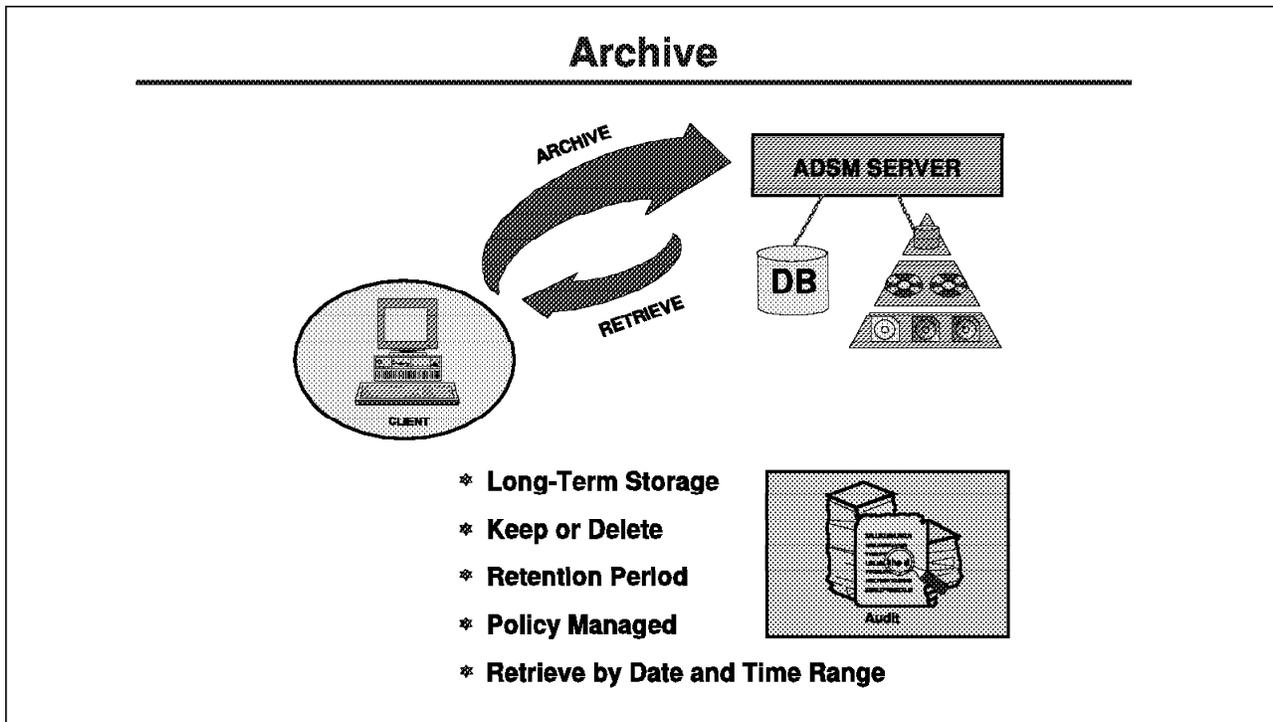
Archive

Archive enables you to save copies of files for some period of time on the ADSM server. Archive copies are never replaced with more current versions (as is the case with backups). They are preserved exactly as you store them.

Long-Term Storage

You use archive when you want to place copies of files in long-term storage. Archived files typically contain data that is unlikely to be referenced in the near future yet must be kept (for example, for legal reasons, vital records, or auditing).

When you archive files, you can associate a description with them, so that at recovery time you can bring back all of the files that match that description.

Keep or Delete

After the files have been archived, you can keep or delete them from the workstation. You could delete the files to reduce the amount of data on your workstation. You can also look at the delete option as a way of reducing your investment in new hard disks for those clients that do not have hierarchical storage management.

Retention Period

Archived copies are kept for a certain retention period and then expired. The retention period can be up to 9,999 days, that is, 27 years and some days.

Policy Managed

Like backups, archived files are managed according to policies. Instead of having rules to keep a specific number of versions, ADSM keeps archived files for a specific time period (for example, seven years).

Retrieve by Date and Time Range

This function works in the same way as described in "Restore by Date and Time Range" on page 15.

Client Function - Hierarchical Storage Management (HSM)

HSM enables you to defer the cost of local disk upgrades as well as the interruption of service when large disks would be required. HSM is also called *space management*.

Migrates Inactive Data

HSM migrates inactive data from the local file system to ADSM server storage. Migration can be controlled by thresholds. For example, when a file system becomes more than 80% full, begin to migrate inactive files. You can also select files to be migrated. Migrated files can be optionally compressed as they are for backup and archive.

Transparent Recall

If a file that is migrated is referenced, it is automatically and transparently recalled. You may experience a slight time delay before they can access the file, but otherwise they are unaware that the file was not in local storage.

Policy Managed

The management of files to be migrated is controlled by policy.

Note: For space management, some controls are in the management class, and some are set by the root user. Include-exclude lists can also be used to control eligibility.

Integrated with Backup

HSM is also integrated with backup. You can specify not to migrate a file until it has a backup. If you migrate a file and then take a backup the next day, ADSM is smart enough to copy the file in ADSM storage and not require a recall to back it up. Not all products in the marketplace today provide this level of integration.

Optional Feature of Selected Servers and Clients

HSM is an optional feature available with most ADSM Version 2 servers. HSM clients are currently supported on AIX 3.2, AIX 4.1, and Sun Solaris 2.5.

For more information about HSM, refer to *ADSM: Using the UNIX Hierarchical Storage Management Clients* (SH26-4030) or the redbook, *Using ADSM HSM* (SG24-4631).

HSM function is also provided for Windows NT and NetWare clients by products from Wang Software. AvailHSM provides HSM for NetWare, and OPEN/stor provides HSM for Windows NT. These products perform HSM for the client and send the migrated files to an ADSM server.

For more information, refer to "AvailHSM for Novell NetWare" on page 182 and "OPEN/stor for Windows NT" on page 184.

Client Features
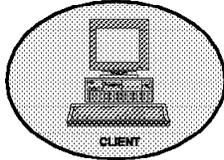
- Graphical User Interface or Command Line Interface
- File Filtering to Reduce Unnecessary Backups
- Restore Options
- Cross-Client Restore or Retrieve
- Restricted Access to Backup or Archive Files
- Data Compression at Client
- Pre- and Post-Schedule Commands
- WebShell Client

© IBM Corporation 1997

Client Features

This graphic lists other client features. For details on client operating systems and the communication methods they support, see Appendix B, "Client Details" on page 283.

GUI or Command Line Interface

The backup-archive client has a GUI and a command line interface (CLI). You can perform more functions from the command line than they can from the GUI.

A GUI is available on all platforms except NetWare. The GUI looks similar to other applications running on Macintosh, Windows, OS/2, and UNIX. The CLI is available on all platforms except Macintosh.

File Filtering to Reduce Unnecessary Backups

File filtering is performed at the client workstation or file server so that files can be excluded from backup.

Each client has a configuration file containing ADSM options. The include-exclude option enables you to specify which files are eligible for ADSM services. The exclude option is used to prevent files from being included from backup processing. The include option is used to make exceptions to the files specified in the exclude option. For example, you might want to exclude the

directory structure where your operating system resides but include the autoexec.bat for backup services.

You can use file filtering to assign a different policy to different types of files.

Restore Options

On restore (and retrieve), you can specify the restore destination, which can be a different location from the location where the files were backed up. When restoring to a location that has a file with the same name, you can also specify whether a file gets overwritten.

You can also restore (and retrieve) files that were backed up within a certain date and/or time range. This would be particularly useful if a virus infected your machine and you want to get back files that were backed up before the virus was introduced. For more information, see "TOTIME and FROMTIME" on page 147.

Cross Client Restore or Retrieve

You can restore (and retrieve) any files you back up. With cross-client restore, you can authorize another client to restore (onto its workstation) files you have backed up or archived.

You can do cross-client restores within an operating system family. For example, within the UNIX family, you can do a cross-client restore between Sun and AIX or between HP-UX and AIX. You can also do a cross-client restore between a high performance file system (HPFS) and another HPFS or between a file allocation table (FAT) file system and another FAT file system. However, you cannot do a cross-client restore between Macintosh and Windows.

Restricted Access to Backup or Archive Files

The default is that only you can restore or retrieve files that you have backed up or archived. You must explicitly grant permission to others before they can restore any files you have backed up or archived.

Data Compression at Client

Data compression is the default and it is done at the client level. Backed up and archived files get compressed before they are sent across the network to the ADSM server. (The compression algorithm used is a variant of Lempel-Ziv-Welch.)

Compression reduces ADSM storage requirements and network traffic. The price you pay for those reductions is an increase in the use of the client's machine cycles to do the compression. (So, for example, if you are using a small DOS machine, you may not want to use compression.)

Compression rates vary, depending on the type of data being compressed. Ratios may range from 2:1 to 3:1. In addition, ADSM's compression algorithm is sensitive to reverse compression. Some files (such as image files or already compressed files) actually grow when you try to compress them. If ADSM is compressing a file and sees that it is growing, it stops and does not compress it.

Pre- and Post-Schedule Commands

Pre-schedule commands specify client operating system commands that you want executed before a schedule begins processing. Post-schedule commands specify client operating system commands that you want executed after the schedule completes.

Pre- and post-schedule commands allow you to perform tasks such as shutting down a database or application before backing up files, and then starting the database or application again after the backup is complete. Another example could be renaming or deleting files after archive.

For UNIX platforms the options are specified in the *dsm.sys* file. On OS/2, NetWare, DOS, and Macintosh platforms, the options are specified in the *dsm.opt* file.

This pre- and post-schedule support is available for all ADSM Version 2 clients except Windows and Windows NT.

WebShell Client

The ADSM WebShell client enables a web browser (on any platform located anywhere on the network) to interface with an ADSM backup-archive client. Through the web browser, you select any of the following client operations through different HTML forms:

- Display status of backup operations

- Incremental and selective backup

- Restore

The selection is converted to an ADSM backup client command passed to the ADSM server. The results of the command are then converted to HTML tags and displayed through your web browser.

The WebShell client is available for AIX, OS/2, and Win32 (Win32 stands for the 32-bit client for Windows NT and Windows 95). For more details, see Chapter 5, "WebShell Client" on page 111.
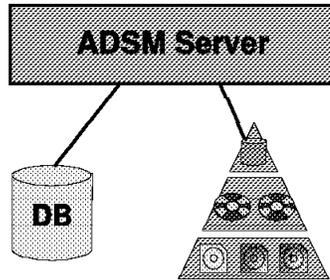
Server Topics

The primary purpose of the ADSM server is to provide a secure environment for the client's backup, archive, and HSM data. Administrator commands are used to manage server operations and storage.

Policy Management

Policy Management

Policy management enables you to define a set of rules explaining how ADSM will treat data.
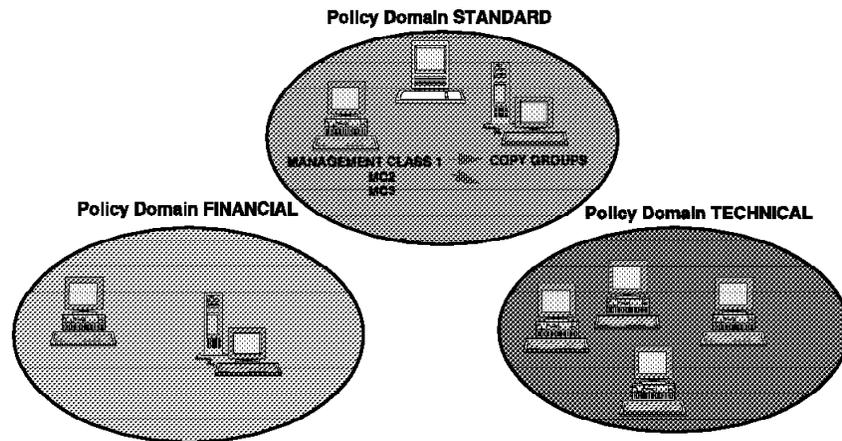
Allows Centrally Defined Policies to Manage Data

Policies are defined centrally on the ADSM server in policy domains. A policy domain is a logical collection of clients. A client belongs to one policy domain. Within a policy domain there are management classes. Management classes define how to manage backed up, archived, or migrated data. (A management class is like the label for a particular policy. There can be multiple management classes within a policy domain.)

Allows Different Types of Data to Be Treated Differently

Policies allow different types of data to be treated differently.

A file is assigned to a particular management class to describe how it should be managed. The flexibility of ADSM enables you to manage your environment as complexly or as simply as you want. To keep things simple, you can have one management class be the default and manage all data in the same way.

Policy Management Parameters
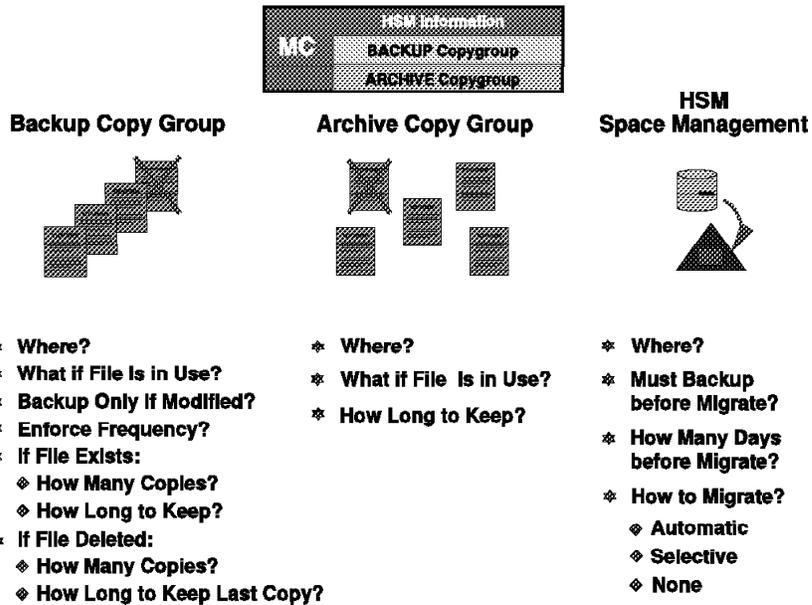
Policy Management Parameters

The parameters for a management class are described in copy groups. There can be up to two copy groups for each management class, one for backup and one for archive.

Backup Copy Group

In the backup copy group, the following is specified:

Where?

Where should the backup copies be stored? They should be stored in a logical destination called a *storage pool*. A storage pool can be disk, tape, or optical. We explain storage pools in more detail in "Storage Pool Assignment" on page 31.

What If File Is in Use?

The administrator tells ADSM what to do if the file is in use while it is being backed up. You can specify four types of serialization:

- Static - If the file is in use, give up. Do not back up the file.
- Shared static - If the file is in use, wait for a while and then retry the backup. The number of retries is set in the client options file, dsm.opt.
- Shared dynamic - If, after a few retries, the file is in use, make a ″fuzzy copy″; that is, read the file and make a copy even if it has changed.
- Dynamic - Do not check serialization, just back up the file.

Backup Only If Modified?

You can specify whether you want a file to be considered for backup according to whether or not it has changed since the last backup. This applies to incremental backups, not selective backups. Two options are available:

- Modified - You want to back up the file only if it has changed.
- Absolute - You want to back up the file regardless of whether it has changed.

For archived files, the mode is always absolute!

Enforce Frequency?

With ADSM you can prevent incremental backups from occurring too often. For example, if the copy frequency is set to 5, five days must pass before ADSM will allow an incremental backup.

If File Exists:

You tell ADSM how many copies to keep of a file while that file exists on the client workstation or file server. The number of copies can be from 1 to 9999! You also specify how long you want to keep the older versions (all except the most recent) until they are expired.

The graphic shows ADSM keeping three versions. When a fourth copy is backed up, the oldest version is deleted.

If File Deleted:

You specify the number of versions to keep if the file has been deleted from the workstation and how long to keep that last remaining copy.

Archive Copy Group

For the archive copy group, there is no "versioning." Each archive copy is its own entity. Therefore, there is much less to specify in the archive copy group.

Where?

The storage pool destination can be the same as or different from the backup copy group.

What If File Is in Use?

A different serialization can be specified for archives.

How Long to Keep?

Archive copies are kept for as long as specified. You can either set the retention period to NOLIMIT, and the file will be retained forever, or you can set a retention period.

The archive copy will expire in *nnn* days from the time the archive was made. The number of days can be up to 9999 days, which is more than 27 years.

HSM Space Management

The space management controls are defined in a management class rather than in a separate copy group.

Where?
The storage pool destination for migrated files can be the same as or different from the backup and archive copy groups.

Must Backup before Migrate?
You can have ADSM wait until a file has been backed up before it is eligible for migration.

How Many Days before Migrate?
You specify the number of days that elapse since the last time the file was referenced before it can be eligible for migration.

How to Migrate?
- Automatic - Allows threshold migration to occur
- Selective - Allows migration by command
- None - No migration is allowed.

You can also specify both automatic and selective.
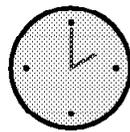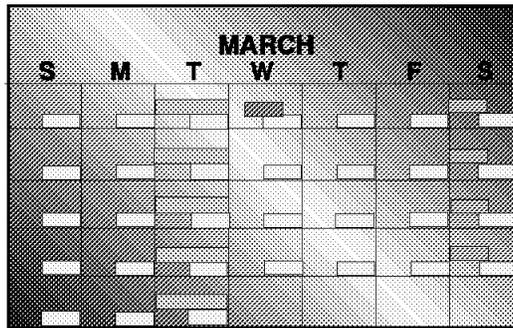
Central Scheduling

Central scheduling enables you to automate the backup and archive process. An administrator is responsible for creating and maintaining the schedules so that a function can be performed on behalf of the client. The administrator associates clients with one or more schedules. The schedules are part of the policy domain and thereby maintained in the ADSM database.

Frequency

The administrator defines the time period (that is, a window) during which a schedule can start and how often to repeat the schedule (for example, every night).

Start and Stop Times

The schedule is run with respect to the server clock. When the client asks at what time to run a schedule, the server specifies x hours. This is important because the clocks on distributed systems, especially workstations, may not have the correct time (because of daylight savings time or other time adjustments).

Target Clients

Once a schedule is defined, it is associated with one or more clients. These are the types of schedules that can be defined:

- Backup-archive

  A schedule can perform a backup or archive on behalf of the client.

- Restore-retrieve

The ADSM V2 server now allows you to define a schedule to perform a restore or retrieve.

- Client OS command

  The ADSM V2 server also allows you to define a schedule to perform a macro or any executable command on that client's operating system. Thus you can do things like send schedule logs or shut down or start up databases on a schedule.

- Administrative commands

  The ADSM V2 server allows you to define schedules to execute administrator commands.

Retry and Randomization

You can specify the maximum number of concurrent clients allowed to log on to ADSM as well as the maximum number of scheduled sessions allowed. If you restrict the number of scheduled sessions allowed on the server, a client is prevented from running a schedule when the maximum number of sessions has been reached. Through the retry option, the client can retry x number of times to run the schedule.

With the retry and randomization options, you have a lot of flexibility to balance the network load.

Let us consider the scenario where you, as an administrator, associate 100 workstations with a backup schedule that has a window of 2 a.m. through 6 a.m. every Friday. You can use the randomization option to tell ADSM to stagger the start times of 100 backup sessions so that they start at different times. In this way, you prevent a big bottleneck, which would occur if all 100 schedules started at 2 a.m.!

Event Log

When schedules occur, they are recorded centrally in the event log at the server. You can view which schedules ran successfully, were missed, and are scheduled to run. You can create an exception reporting list. In this way you can view only those schedules that failed.

Detailed results of the schedule are logged at the client level. High-level results are sent to the server to an event log, so you can view whether or not schedules ran successfully.

You determine how long a record is retained in the log.

Authentication

ADSM is a client/server application that operates over a network. In this environment, someone can "masquerade" as either a client or a server. It is important that the authentication be a two-way process where both client and server verify that they are who they claim to be.

Mutually Suspicious

ADSM's authentication is mutually suspicious. The client is suspicious of the server and the server is suspicious of the client until the authentication process is complete.

Password Not Sent across Network

During the authentication process, the password is not sent across the network.

The authentication process is described as a "conversation."

1. The client begins the conversation by saying, "I am client X; if you have the same password that I have, you will understand this message." (The client uses the password as an encryption key for the message.)
2. The server then receives the message and responds, "I am server Y; if you have the same password that I have, you will understand this different message." (The server uses the client's password as stored in the ADSM database to decrypt the client's message and to encrypt the server's message.)
3. Client X acknowledges that server Y is Ok.
4. Server Y acknowledges that client X is Ok.

This authentication scheme is Kerberos-like, where the trusted third party is the password.

Storage Pool Assignment
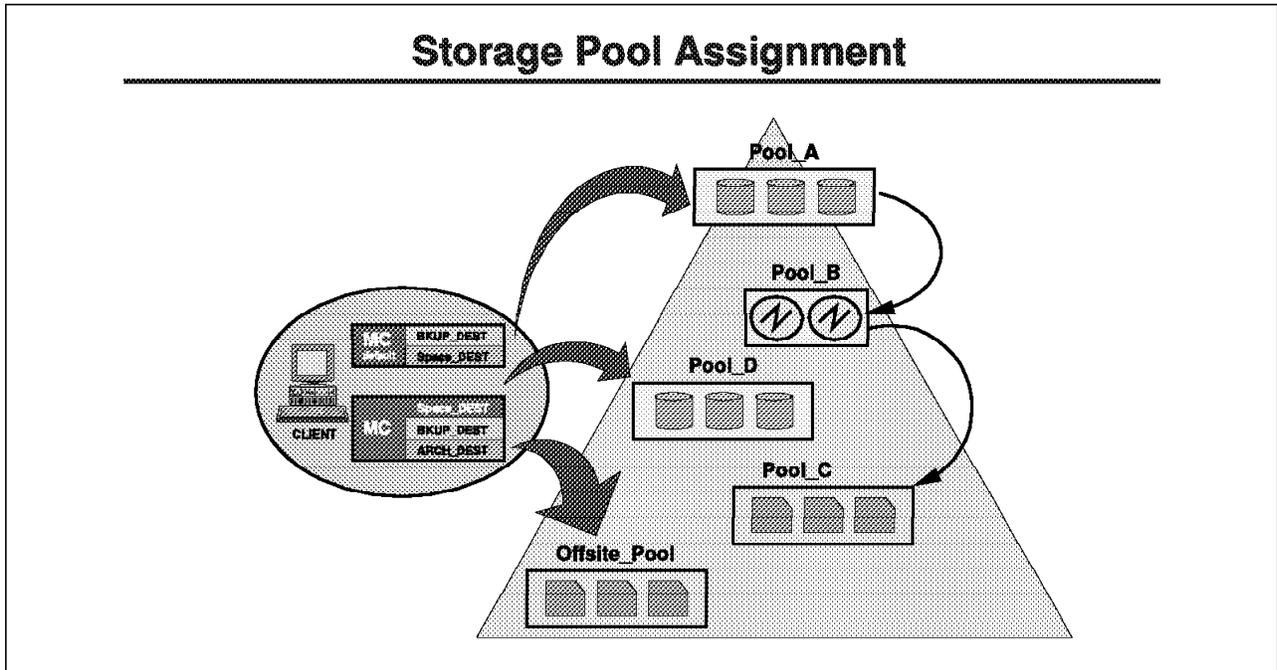
Storage Pool Assignment

ADSM is flexible in how it defines and controls the storage hierarchy. You can have an infinite number of levels in a hierarchy, each called a storage pool.

A storage pool can be disk, tape, or optical. Backups, archives, or migrated files can start on disk or go directly to tape. The management class determines where the client data enters the storage hierarchy.

ADSM controls the movement from one storage pool to the next according to the rules that you define in the storage pool.

A storage pool consists of "volumes." For tape, they are actually physical tape volumes. For disks, they may be physical or logical volumes. For example, on ADSM for MVS, the storage pool volumes are actually VSAM linear data sets. On ADSM for AIX, the volumes can be either a logical volume or a journaled file system file.

Not only is the ADSM storage hierarchy flexible, it is also dynamic. You can add or remove volumes without interrupting server operations. For example, if you install a new type of device as a disk pool, the ADSM administrator can move the storage pool data from the old pool to the new pool without shutting down the server. Or, if you have to add space to a storage pool, you can easily define new volumes and thereby expand the size of the storage pool without disrupting service.

For more information about the specific types of devices that ADSM supports, refer to Appendix C, "Device Support" on page 323.

Storage Pool
Hierarchy Management and Space Reclamation
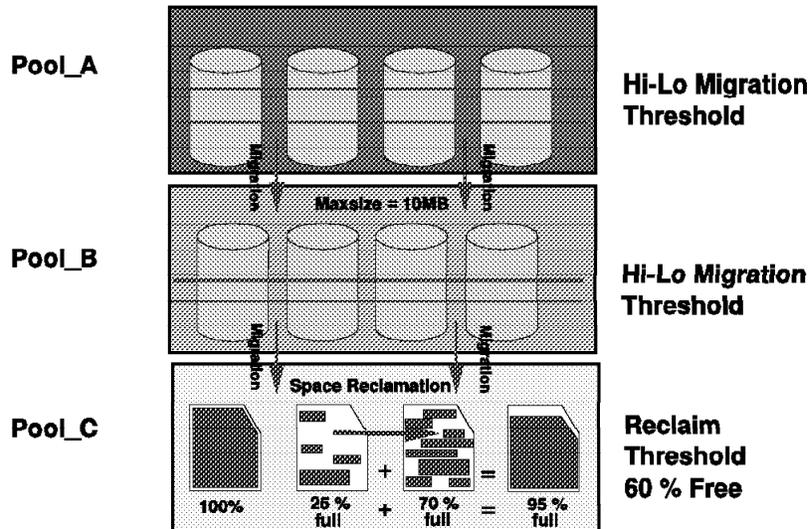
Storage Pool Hierarchy Management and Space Reclamation

Three controls are available to help you automatically control the space in the storage pools:

- Migration
- Maxsize
- Space reclamation

Migration

Migration helps you control the amount of free space within a storage pool. You can define a high-low (Hi-Lo) migration threshold for each storage pool in the storage hierarchy. These thresholds tell ADSM when to move data from one storage pool to another. On this graphic, data will be migrated from storage pool POOL_A to storage pool POOL-B. If there is not enough space in POOL_B, migration of data in POOL_A will wait until migration on POOL_B has freed enough storage (by migrating data from POOL_B to POOL_C).

How does migration work? When the amount of data in a storage pool reaches the high threshold, ADSM moves client data to the next storage pool, until the storage pool reaches the low threshold. ADSM's objective is to clear as much space as quickly as possible. To do this, ADSM picks the client node whose collection of data occupies the most space in that storage pool and migrates that client's data to the next storage pool. ADSM then picks the next client node that has the most data and migrates it. ADSM continues the migration until it reaches the low threshold. This operation not only clears space quickly but also keeps a client's data close together in the hierarchy.

Maxsize=10MB

If the administrator assigns a maxsize to a storage pool, files can be placed in that storage pool provided they are equal to or less than the assigned maxsize. If compression is used, the uncompressed size of the file is used in the comparison. On this graphic, POOL_B has a maxsize of 10 MB. If a file is to be backed up into POOL_B and its size is greater than 10 MB, POOL_B will be bypassed, and the data will be directly written to POOL_C.

Maxsize is not required. If you do not assign a maxsize, and a file entering the storage pool causes it to exceed its high threshold, migration will occur to move data to the next storage pool.

Space Reclamation

Space reclamation helps you control the number of tapes needed to support a storage pool and ensures that the tapes are being used as efficiently as possible. Reclamation applies only to tape storage pools.

Initially, tapes are typically 100% utilized. However, over time data is typically deleted or expired from the tape, and there is then unused space on the tape. With ADSM, you can specify a reclamation threshold. When the percentage of "empty" space on a tape exceeds the reclamation threshold, ADSM will find another tape to which to move the data. The now empty tape can be reused or recycled.

In this example, reclamation will begin when the amount of reclaimable space on a tape exceeds 60%. The tape that is 25% full has 75% reclaimable space, so it is eligible for reclamation. The valid data on tape will be copied onto another tape in the storage pool with the closest fit, in this example, a tape that is 70% full.

Storage Pool - Collocation

In most cases it is necessary to keep the user data of one client on one volume of sequential media. This function is provided with collocation.
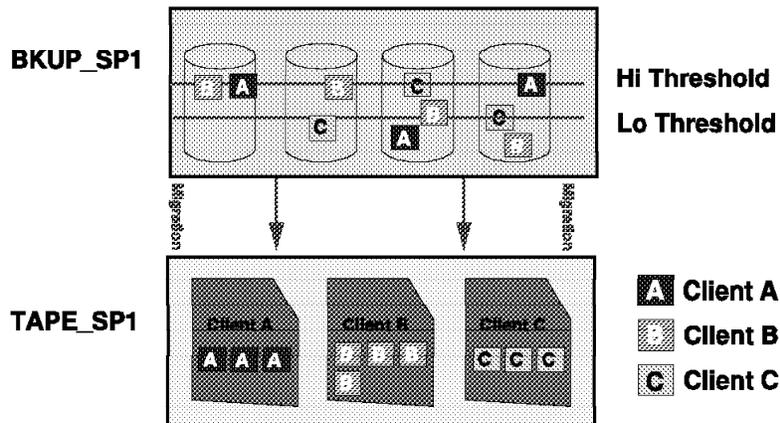
Keep Workstation Data Together

Data that is migrated to tape, over a period of time, can be stored among various tapes. Collocation directs the ADSM server to keep the data from one client on a minimum number of tapes. You can choose to have the ADSM server collocate at the client or filespace level (see "COLLOCATION by FILESPACE" on page 166). A filespace is an entire drive or file system belonging to a client.

Reduce Tape Mounts during Restore

The benefit of using collocation is that you reduce the time it can take to restore a large group of files.

Selectable Option at Storage Pool Level

You can collocate a client on individual filespaces. Collocation is allowed for sequential storage pools (tape, optical).

On this graphic, with collocation on, client A's data is kept on client A's tape, client B's data is kept on client B's tape, and client C's data is kept on client C's tape. Collocation applies whether you are migrating data from a disk to a tape storage pool or going straight to tape with no disk storage pool.

Copy Storage Pools
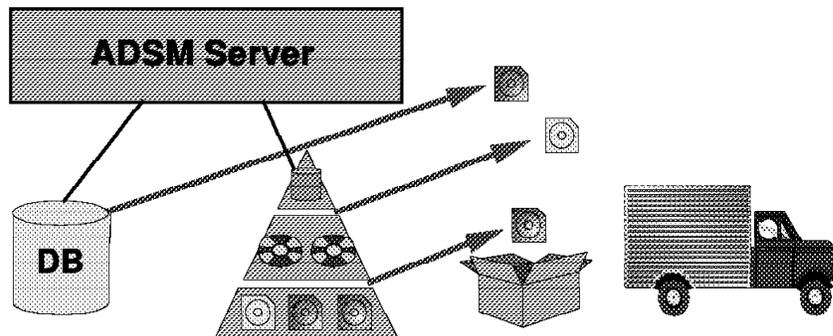
- Offsite Backup of Database, History Files, and Storage Pools
- Available on Selected Server Platforms
- EXPORT Provides Offsite Backup in Earlier Versions

Copy Storage Pools

For disaster recovery purposes it is necessary to produce an offsite backup of the ADSM database and storage pools.

Offsite Backup of Database, History Files, and Storage Pools

Two types of ADSM server database backups are provided, full and incremental. You can run full or incremental backups as often as necessary to ensure that an ADSM database can be restored to an acceptable point in time.

A new type of storage pool, *copy storage pool*, is provided in ADSM Version 2. Copy storage pools are sequential storage pools (including DEVTYPE=FILE) that are used to store duplicate copies of files residing in any type of primary storage pool: disk, optical, or sequential media. A client cannot directly store files in a copy storage pool.

Copy storage pools are provided with offsite volumes for disaster recovery. ADSM never automatically uses offsite volumes in the disaster recovery (offsite) storage pool. The volumes in this copy storage pool should be marked with the new access mode, OFFSITE, to ensure that ADSM does not request a volume mount.

Store the offsite volumes in the disaster recovery copy storage pool offsite together with the database and history file backups to ensure that you can recover your ADSM server in case of a disaster.

Available on Selected Server Platforms

Offsite backup of the ADSM database, history files, and storage pools is available on all ADSM Version 2 server platforms and can be performed when the server is up and running, online!

EXPORT Provides Offsite Backup in Earlier Versions

In ADSM Version 1, you have to shut down the ADSM server to produce offsite backups of your database and storage pools to get a consistent backup.

To back up the database, use your favorite backup utility program. It is the fastest way to get an offsite backup of the database.

You can use EXPORT to copy storage pool data to sequential media. Using this feature you can back up data in storage pools and store the sequential media at an offsite location.
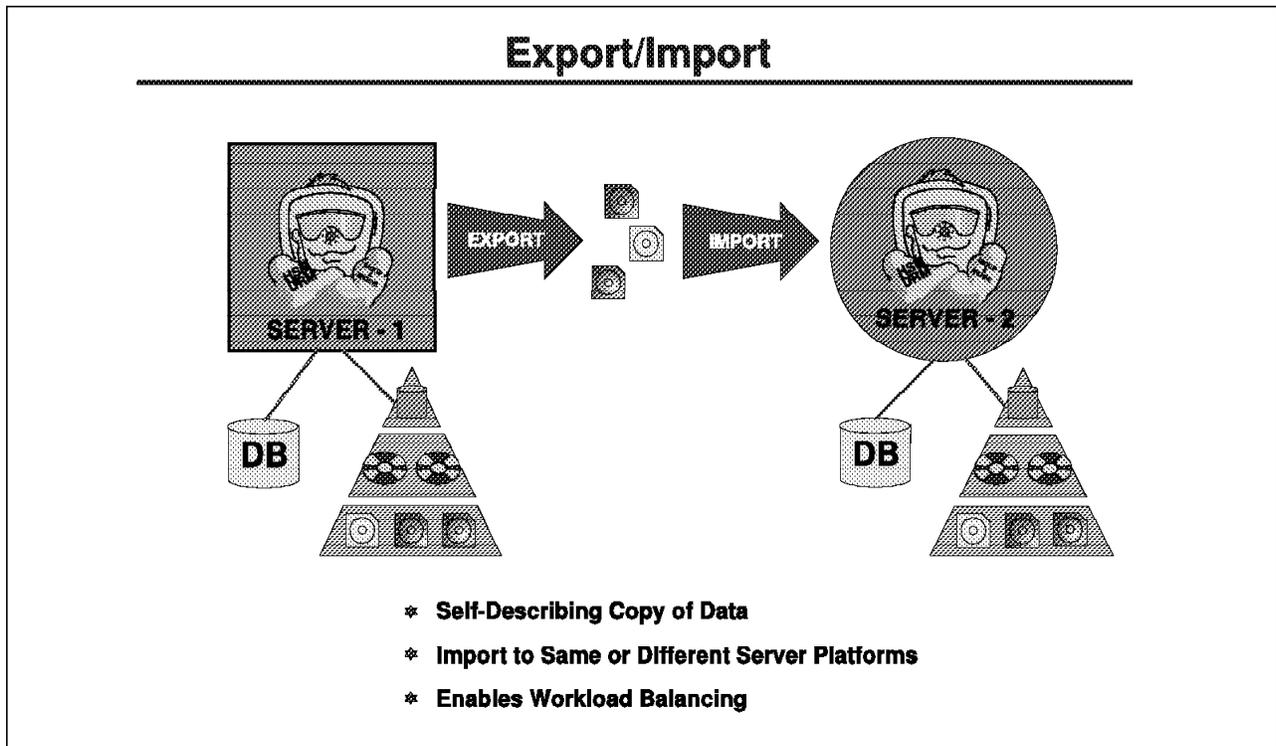
Export/Import

Export/import enables you to move ADSM data from one server to another. You can export or import administrator, node, server, and policy information.

Self-Describing Copy of Data

The export commands create a self-describing copy of specified server information. The output media can be a tape or a flat file. You can export database information as well as the backed up and archived files stored in the storage pools.

Import to Same or Different Server Platform

You can import to the same or a different server platform. You just have to ensure that the server to which you are importing supports the same tape format as the server from which you exported, if you are exporting to tape media.

Enables Workload Balancing

Export/import facilitates workload balancing by allowing you to move the backups and archives associated with a client to another server. Additionally you can use the export/import commands to clone information, such as policies, from one server to another.

ADSM Database

```
                         ADSM Database
```

write - to log and DBs   read - from either DB

DB & Mirror(s)   Log
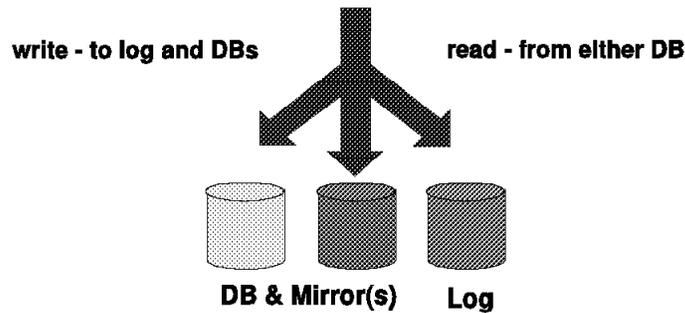
* **Database**
  * **Client Node Information**
  * **Client File Information**
  * **Policies**
  * **Schedules**
* **Recovery Log**
  * **Describes Changes to DB**

* **Mirrored for High Availability**
* **Automatic Alternate Volume Switching**
* **Dynamically Expand or Contract**
* **Full or Incremental Backup while Server Active**

ADSM Database

The ADSM database is the heart of the server. It is highly tuned for storing and retrieving data from the storage pools. It is a true database except that it does not have a general query language and does not allow users to create tables.

There is no client or administrator interface directly into the ADSM database. Transactions such as defining a node, performing a backup, and changing a management class will update the database. The administrator can perform queries from the CLI or view information from the GUI that will read the database.

Database

The database keeps track of where all backed up and archived data is located in the storage pools. All policy and schedule information is stored in the database. Because the location and attributes of a file are stored in the database, you can easily view the attributes and versions of a file when you have to restore it, without reading and searching through all of your tapes to find the right one.

The database has all of the features associated with a database management system such as:

- Transaction orientation
- Write-ahead logging
- Transaction commit, rollback, and redo
- Point-of-failure recovery
- Hierarchical intent locks

Recovery Log

The recovery log is used to help maintain the integrity of the database. It keeps track of all changes made to the database, so that if a system outage were to occur, a record of the changes would be available in the log.

When a change to the database occurs, the recovery log is updated with some transaction information before the database is updated. Thus uncommitted transactions can be rolled back during recovery so that the database remains consistent.

Mirrored for High Availability

Using mirroring, you can configure ADSM to maintain as many as three copies of the recovery log, the database, or both on independent storage volumes. The mirrored copies are treated equally; there is no concept of primary copy and alternate copies. Therefore the server reads from the database with the best response time!

Automatic Alternate Volume Switching

If a physical volume that the database or recovery log resides on fails, ADSM takes the failing volume offline and continues to perform operations using the remaining intact copy or copies. ADSM service is not interrupted.

After the failed volume is repaired, ADSM allows the volume to be brought online and synchronized with the intact volumes, again without interrupting service.

Dynamically Expand or Contract

Space can be dynamically added or deleted from the database or recovery log as needed. For example, you can start with one volume. As the system grows, you can allocate more space to the database or recovery log without having to shut down the server.

Full or Incremental Backup while Server Active

With an ADSM V2 server, you can perform full or incremental backups of the database. There are trade-offs between running full or incremental backups:

- A full backup takes longer to run than an incremental backup because it copies the entire database. However, recovery time is faster with a full backup because only one set of volumes has to be loaded to restore the entire database.

- Incremental backup takes less time to run because it copies only those database pages that have changed since the last time the database was backed up. However, incremental backups increase the time it takes to recover a database because a full backup must be loaded first, followed by some or all of the incremental backups in the same database backup series.

  You must run a full backup for every 32 incremental backups. The average is to perform incremental database backups daily, and a full backup weekly.

Disaster Recovery Manager (DRM)

With DRM you can better prepare and execute a disaster recovery plan.
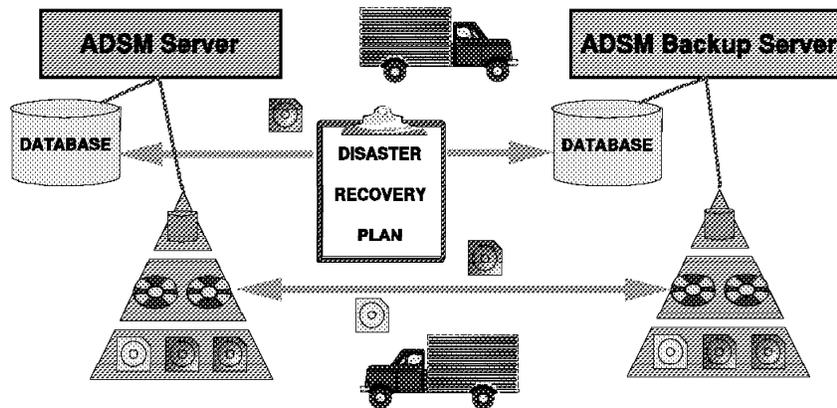
Automatic Generation of ADSM Server Recovery Plan

With the PREPARE command you can automatically query the required information from the ADSM server. The resultant information is stored in the recovery plan file.

Determine Which Media to Move Offsite, Onsite

Knowing the location of offsite and onsite recovery media is critical to successful disaster recovery. With DRM you can determine which volumes have to be moved from offsite to an onsite location for recovery purposes.

Inventory Machines in ADSM Environment

DRM enables you to store client machine information, such as machine location and characteristics, in the ADSM database.

Use the stored information to guide you through disaster recovery.

Optional Feature of Selected Server Environments

DRM is an optional feature currently available for the MVS and AIX Version 2 servers.

For more information refer to Chapter 2, "Disaster Recovery Manager" on page 73.
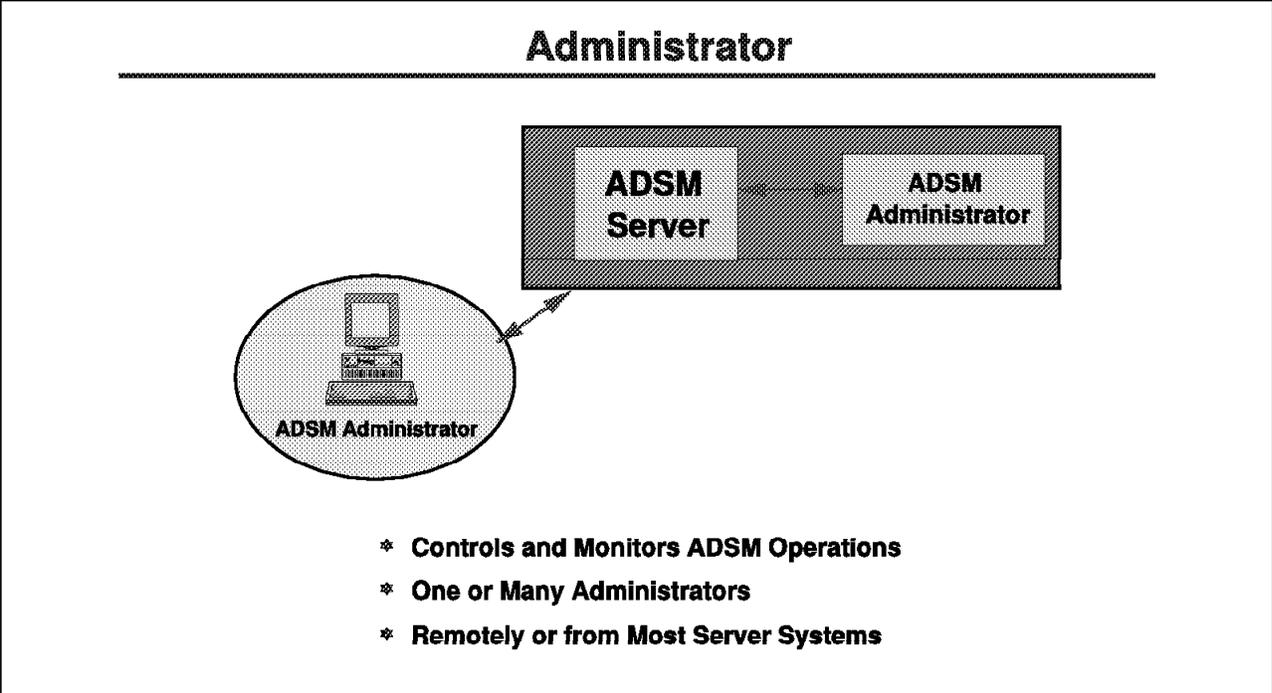
Administrator

Administration can be done from either the server console or an administrative client. We look at some of the basic tasks of an administrator.

Controls and Monitors ADSM Operations

The ADSM administrator tailors and controls ADSM operations. Administration can be done from either the server console or an administrative client. The administrative client can run on either a mainframe or workstation platform and, depending on the platform, can be a CLI or GUI.

One or Many Administrators

One administrator or multiple administrators can manage the ADSM environment.

Remotely or from Most Server Systems

An ADSM administrator can be located anywhere in the network and on a different platform from your ADSM server.

```
┌─────────────────────────────────────────────────────────────────────┐
│                      Administrator Features                          │
│  ▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬ │
│                                                                       │
│                                                                       │
│              ❋ GUI on OS/2, AIX, HP, Sun, Win32                       │
│                                                                       │
│              ❋ CLI on Other Clients                                   │
│                                                                       │
│              ❋ Multiple Levels of Authority                           │
│                                                                       │
│                   ◈ System                                            │
│                                                                       │
│                   ◈ Policy                                            │
│                                                                       │
│                   ◈ Storage                                           │
│                                                                       │
│                   ◈ Operator                                          │
│                                                                       │
│                   ◈ Analyst                                           │
│                                                                       │
│             ❋ Centralized or Decentralized Control                    │
│                                                                       │
└─────────────────────────────────────────────────────────────────────┘
```

Administrator Features

There are some administrator features that help you to manage your ADSM environment.

GUI on OS/2, AIX, HP, Sun, Win32

You have full administrative GUI support on the client types listed.

CLI on Other Clients

You can also choose to work with the CLI, which is available for most of the client platforms. For specifics see "Miscellaneous ADSM Client Enhancements" on page 139.

Multiple Levels of Authority

Different levels of authority can be defined for the administrator. The levels are:

- System - all-powerful, can define other administrators
- Policy - manage policies, define management classes, and register nodes
- Storage - add volumes, define storage pools, manage database and recovery log
- Operator - respond to system messages, cancel sessions
- Analyst - use query commands only

Centralized or Decentralized Control

ADSM can be either centrally or decentrally controlled. The administrator's scope of authority can be centralized or distributed, as you can see on the next graphic.

Administrator Span of Control
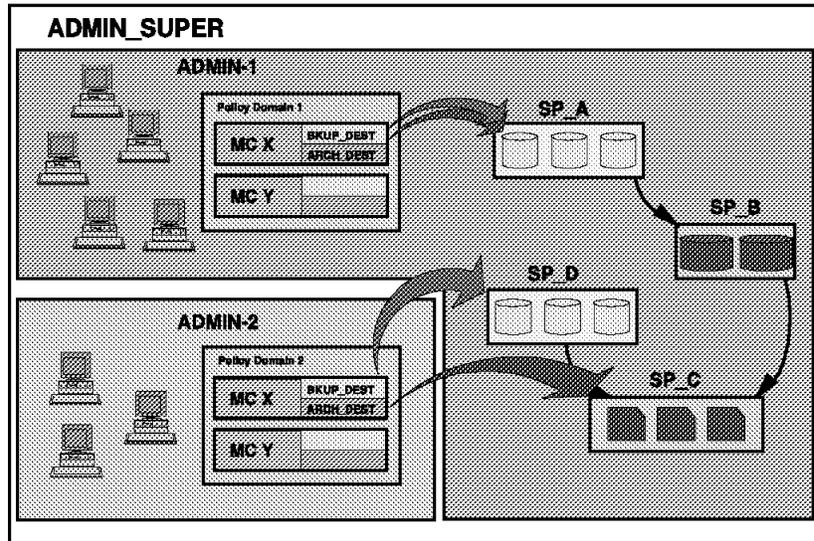
Administrator Span of Control

Some customers may want their LAN administrators to have control over their domain. You can set it up so that LAN administrators have policy authority, which allows them to add users, change passwords, and define the ADSM policies. You can restrict the domain over which they have this policy authority.

On this graphic:

- ADMIN-SUPER has system authority.
- ADMIN-1 has policy and storage authority over policy domain 1 and all storage pools.
- ADMIN-2 has policy authority only over policy domain 2.
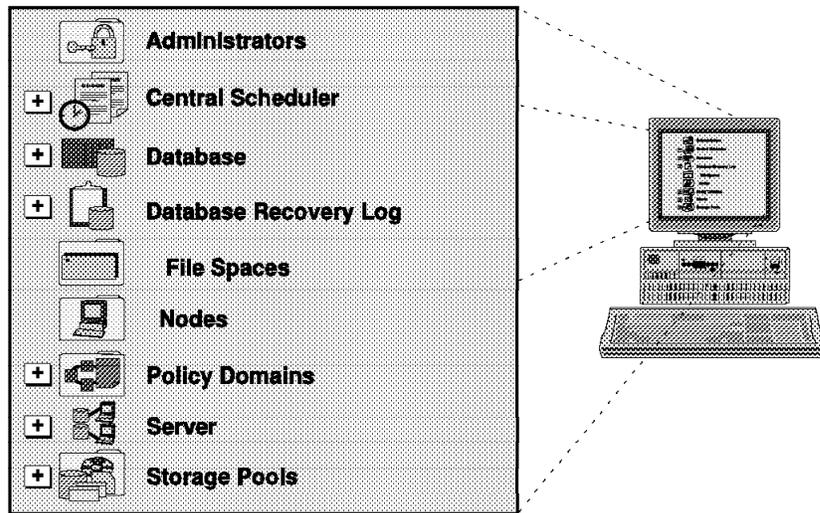
Administrative GUI

This is the main menu of the ADSM Version 2 administrative GUI. Each icon represents a different class of object.

- Tree structure (classview)

  With classview, it is now possible to have a tree view of ADSM icons by simply clicking on the boxed ″+″ symbol to the left of the icons.

- Fast path selection

  With this tree structure you can directly select any of the icons on the branches. You can directly select certain ADSM administrative functions without having to traverse through several administrative windows. This is called fast path selection.

This GUI works with an ADSM V1 server as well as an ADSM V2 server.

Supported Storage Devices
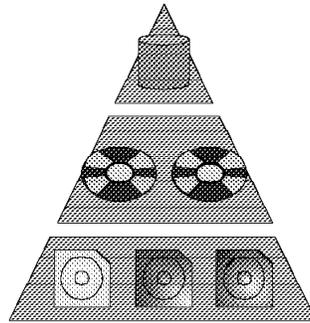
This graphic shows some of the supported storage devices.

IBM and Non-IBM Storage Devices

ADSM supports a wide range of IBM and non-IBM hardware technologies.

Disk, Optical, Tape, Jukeboxes, Libraries

The hardware supported includes disk, optical, tape, jukeboxes, and libraries.

Attachment Varies by Platform

Different ADSM servers support different storage devices.

For more information, refer to Appendix C, "Device Support" on page 323.

Other Topics
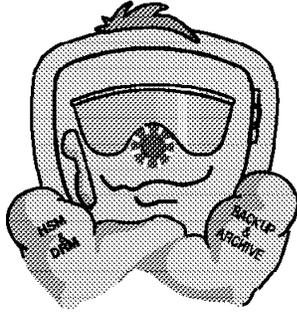
* Common Backup Paradigm

* Another Backup Paradigm

* ADSM Backup Paradigm

* Network Bandwidth Considerations

* Database Support

* ADSM and Complementary Products

* ADSM and Tivoli

* ADSM Information Sources

* ADSM Redbooks

* Implementation Services

* IBM 3466 Network Storage Manager

Other Topics

Common Backup Paradigm - 1

Full Backup	Weekend
Incremental	Monday
Incremental	Tuesday
•	Wednesday
•	Thursday
•	Friday
Full Backup	Weekend
Incremental	Monday
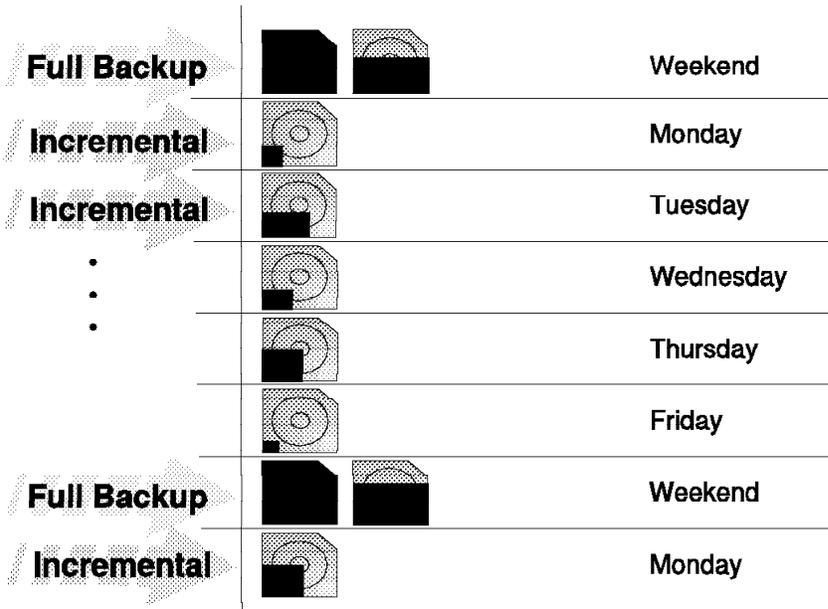
Common Backup Paradigm - 1

A key benefit of ADSM's approach to backup and restore is that restoring a single file is a fast and easy process. You can pick the file that you need to restore from an online catalog (stored in the ADSM database) even when the backup copies are on tape.

However, there are times when you need to restore an entire system. In addition, many people are confused when they learn that ADSM does not do full backups. This section explains why you do not have to do a full backup with ADSM and demonstrates how ADSM's paradigm saves you time during backup and restore.

In today's networking enterprise, the most common way of performing backups is to take full backups on a regular basis, with incremental backups between each full. A full backup backs up every file in the file system or drive. An incremental backup backs up files that have changed since the previous incremental backup (or the full backup if no previous incrementals exist).

To avoid the management of too many tapes, the number of incrementals should be as few as possible. The average frequency is one full backup every week, plus five or six incremental backups (one per day) in between.

This way of performing backups implies:

 • One tape per day

 • Very little data on each tape (except the full backup tapes)

 • When performing the second full backup, you ignore all of the previous full backups, erase the tapes, and send them back to the scratch pool.

The administration of the tapes, inventory and tracking, tape labeling, and archiving must be done manually in most cases.

In addition, each time you do a full backup (on the weekend), you send *all of the data* again. This is neither a practical nor time-saving process.

Common Backup Paradigm - 2

Restore on Friday morning = a multiple-step restore operation

Restore the last Full Backup

+ Restore Monday Incremental

+ Restore Tuesday Incremental

+ Restore Wednesday Incremental

+ Restore Thursday Incremental
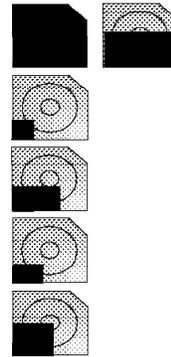
Common Backup Paradigm - 2

Let us try to restore now! We have one full backup, plus several incrementals, each containing a little data. A restore operation on Friday morning involves the restore of the last full backup, plus the restore of each incremental backup that was performed between the weekend and Friday. How many tapes will be involved overall? How many tapes for each client? Which tape belongs to which client?

Another Backup Paradigm - 1
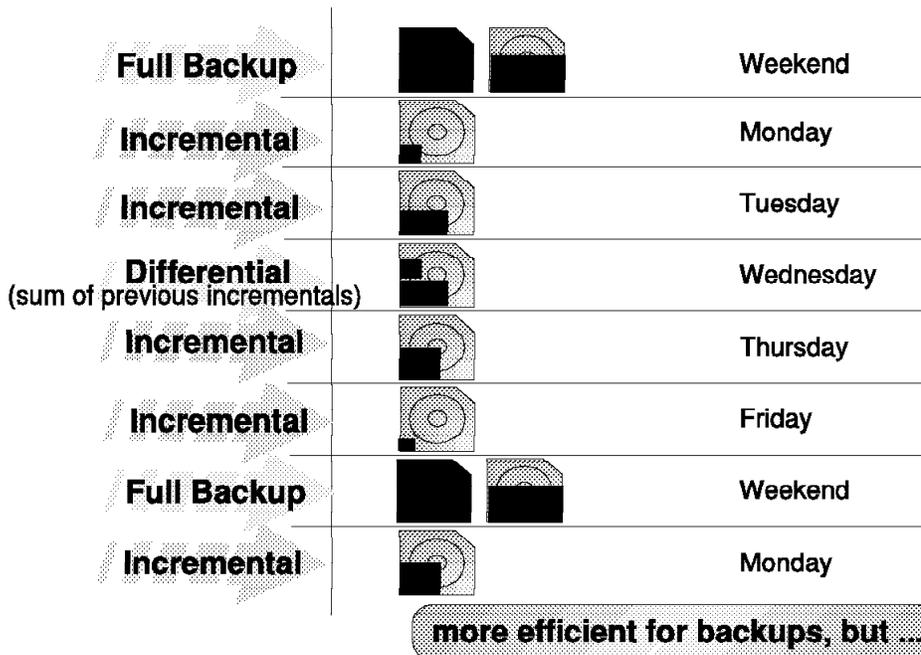
Another way of performing backups is to take full backups and differential backups with incremental backups in between. In this paradigm, a full backup backs up every file in the file system or drive; a differential backup backs up the files that have changed since the previous full backup; and an incremental backup backs up files that have changed since the previous incremental backup (or the previous differential backup if no previous incrementals exist, or the previous full if no previous differentials exist).

This process reduces the number of tapes to manage because you can discard your incremental tapes once you have done a differential. But you still have to manage the incremental tapes.

This way of performing backups implies:

- One tape per day
- Very little data on each tape (except the full backup tape)
- More tapes to manage, because you have to keep the full backup tapes, the differential tapes, and the incremental tapes.

Using a such backup procedure is an improvement over the full and incremental paradigm, but restore is still a multiple step operation, and you must still take periodic full backups—which backs up all of your files again regardless of when they've been changed.

Another Backup Paradigm - 2
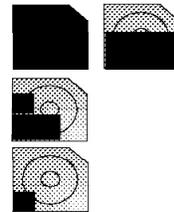
When restoring data that has been backed up with the full plus differential process, the complexity of the restore solution is a little simpler, but you still have to restore from several tapes, it is a multiple step process, and it is time-consuming.

The ADSM Backup Paradigm - 1

The ADSM Backup Paradigm - 1

Base Backup		Weekend
Incremental		Monday
Incremental		Tuesday
Incremental		Wednesday
Automatically Reclaim Empty Space		
Incremental		Thursday
Incremental		Friday
Incremental		Weekend
Backup is faster, and ...	Automated Tape Reclamation and Collocation Processes	
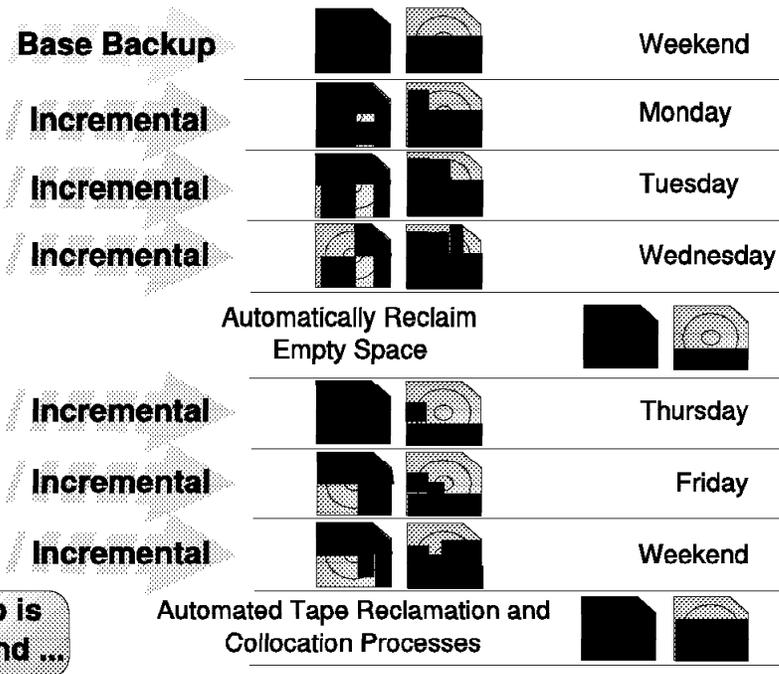
The ADSM Backup Paradigm - 1

The most critical factor in any backup is the amount of data to be moved. If you agree that "the fastest way of moving data is not to move it," ADSM's incremental backup is the best way of backing up files because you back up a version of a file only once. It takes a little more time up front for ADSM to figure out which files need to be backed up and which do not, but you do not have to move every single file to the backup location every single day.

Now let us look how ADSM would do this job for you:

- With ADSM you only have to do incremental backups.

- The very first incremental backup is a "full" backup. You *NEVER* need another full backup.

- All other incremental backups back up only the files that have changed since the previous incremental.

- The amount of data sent across the network is dramatically reduced.

- The collocation option places the client data from each backup on the same tape. You only have to use another tape if the first tape becomes full.

  If you have clients with extremely large file systems, you can even collocate file systems separately.

- At the same time, a regular (usually daily), automated expiration process removes from the ADSM database all entries related to expired files (for example, old versions, expired retention periods, and intentionally deleted backups). This process creates empty spaces on tapes. These spaces are

then handled by the reclamation process, which keeps the tapes as few and as filled as possible. If the free space of a tape reaches the defined threshold, two or more tapes are copied onto one.

- The administration of the tape and its contents is done automatically. The ADSM database stores all related information.

A tape library, such as IBM's 3494 or 3495 Tape Library Dataserver, used in conjunction with ADSM is very beneficial. It automates tape handling and reduces the wait time for tape mounts.

The ADSM Backup Paradigm - 2

ADSM was born to restore! When restoring data that has previously been backed up by ADSM, you will notice that:

- Most of the time, no more than one or two tapes is mounted for each client, because tapes are filled.

- Each client has its data on its own tapes.

- In the case of a full restore, the process is much faster.
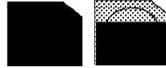
The permanent incremental backups and the asynchronous data consolidation provided by the ADSM server through collocation and reclamation are the basic principles of the ADSM backup paradigm.

No Unnecessary Backups

By backing up only changed files, ADSM saves time, disk space, tapes, and network usage. The backup windows are reduced. The backup/archive functions of ADSM are thus fully in accord with ADSM's storage management functions. ADSM not only protects the data but also manages the storage environment efficiently.

Collocation

Because of the collocation process, even after months of incremental backups, ADSM requests only the few tapes that contain your client's data. Your client's data is not spread all over your tape library. The ADSM database keeps track of which files are on which tapes. Therefore, ADSM does not waste time in READ/SEARCH operations looking for the files to restore. The restoration process is faster and automated.

Reclamation

The reclamation process helps fill the tapes more efficiently, avoiding an endless growth of your tape library. This is valuable to customers who manage tapes manually as well as customers who use automated tape libraries.

Network Bandwidth Considerations

The subject of performance has always been a complex area, even on a single platform. That is why the common answer to any performance question is: "It depends."

A critical point of performance is the network bandwidth.
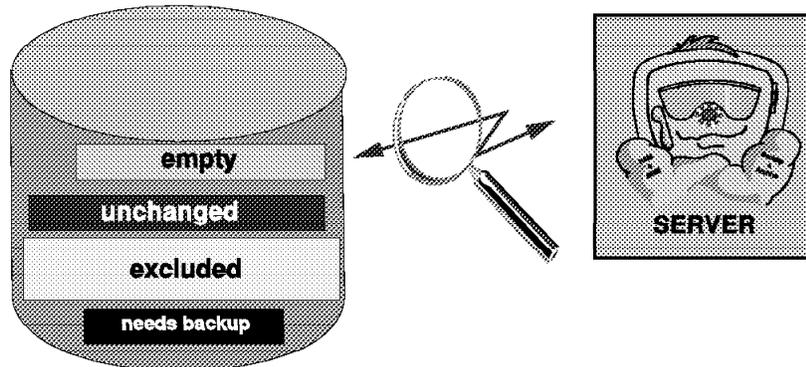
Only a Small Percentage of User Data Changed Daily

As we mentioned earlier, with ADSM you should use only incremental backups because only a small percentage of user data changes daily.

Redundant "Full" Backups Unnecessary

Periodic full backups in ADSM accomplish nothing; restores do not run faster after a full backup than after multiple incremental backups, because either way the ADSM database contains the same information and the storage pool has the same file images. If the file has not changed, backing it up again does nothing but waste resources!

Scheduling during Idle Times

Find out when your network traffic and server CPU usage are low (usually at night). Schedule the backups at that time and choose the randomization option for the server scheduling.

Compression

ADSM gives you the option of compressing the data on the ADSM client machine before sending it over the network. Compression has some clear advantages:

- Less data is sent over the network.
- Less disk or tape space is used for the storage pool.
- Fewer cycles are spent on the server handling the data.

Network Capacity

LANs are rated in megabits per second, and most system throughput values are stated in megabytes per second. The factor of 8 bits per byte is easy to overlook. Even after converting network speeds to megabytes per second, the number represents the maximum speed of the network when data is flowing. Remember that the communication protocol adds headers, control data, and acknowledgments to the data sent through the network. A rule-of-thumb is to use 50%-60% of the published capacity as the practical limit of backup throughput on a token ring, fiber distributed data interface (FDDI), or asynchronous transfer mode (ATM) network, and 35% on an Ethernet.

Network Tuning

Many LANs are geared for the transmission of small messages such as are found in mailing applications. To avoid inefficiencies ADSM needs a large message packet to handle the large volume of bulk data. Set network tuning parameters to provide large packet sizes and substantial buffer sizes. Even when these parameters are set optimally, other obstacles can limit performance. An old bridge or gateway can force a reduction from a large packet size to a smaller size, which would dramatically reduce throughput.

Database Support
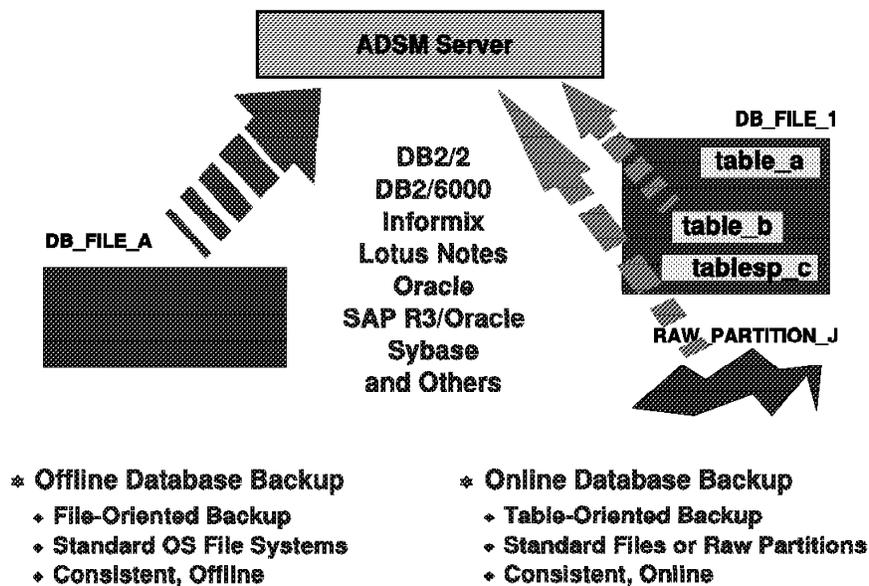
Database Support

Here are some examples of applications that use the ADSM API.

Offline Database Backup

Today, with any version of ADSM, you can back up any database or application that stores its data on files. You will probably have to stop the database manager or application and make sure that all buffers have been flushed to disk. You can then use ADSM to back up the files that the database or application uses.

If you are using a UNIX database that stores its data in raw files, you can still use ADSM to back it up. However, you must first use a database utility (usually supplied with the database) to dump the database from a raw device to an ordinary file. You can then use ADSM to back up those files.

Online Database Backup

When a product uses the ADSM API, it has the potential to provide:

- Online backups - The product provides a consistent backup of the data without stopping the database or application.

- Incremental backups - The product has knowledge of the internals of the data and can therefore back up only the changed portions.

- Raw device support - The product has knowledge of how the data is stored and can back up the data from the raw device without creating an intermediate file.

Implementation of online backups, incremental backups, and raw device support depends on the program using the ADSM API. Some products that use the ADSM API are:

- DB2/6000 - Versions 1 and 2 use the ADSM API to send backups to any ADSM server.

- SQL-BackTrack - Runs database administration tasks on Oracle or Sybase databases.

- Lotus Notes - ADSM is the first product to provide incremental backup of Lotus Notes databases.

For detailed information refer to Chapter 12, "ADSM and Databases" on page 201.
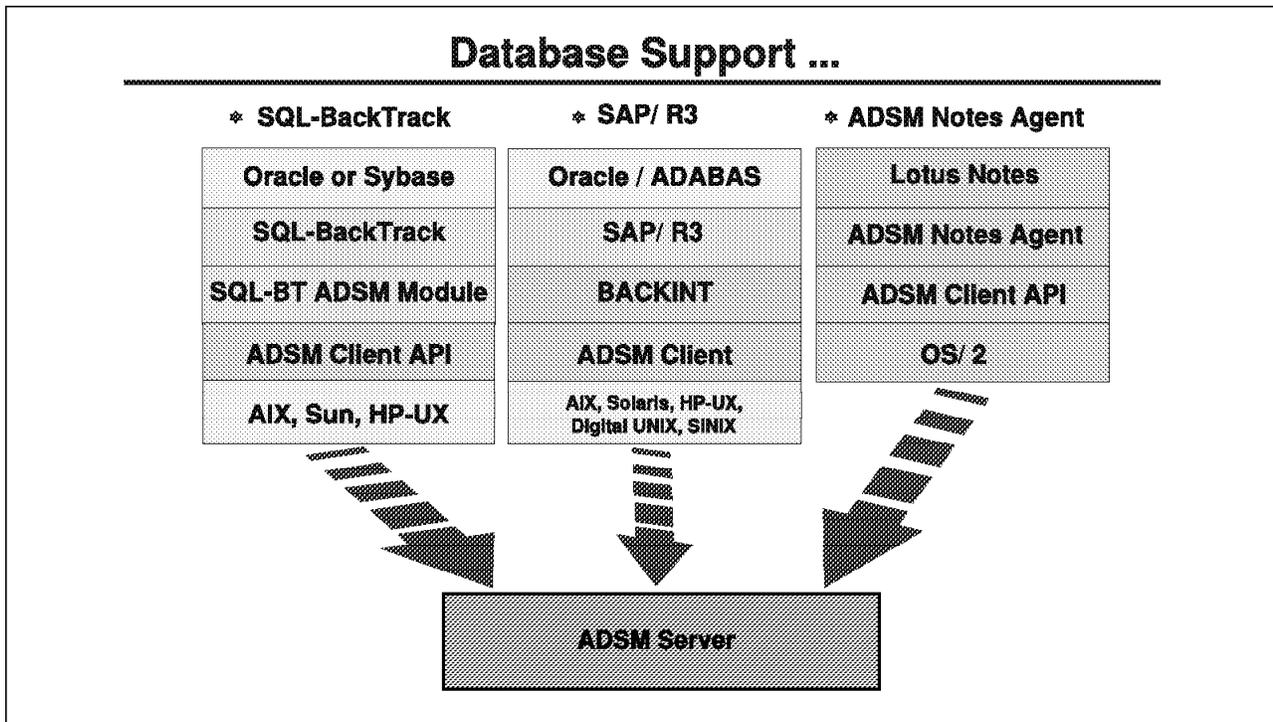
Database Support ...

Here are some examples of products that use the ADSM API.

SQL-BackTrack

SQL-BackTrack from DataTools, Inc. offers additional functions and more flexibility for backing up Oracle and Sybase databases. It provides online database backup, raw device support, and incremental backup of databases down to the 2K page level.

An optional SQL-BackTrack ADSM module is available from DataTools that integrates SQL-BackTrack database backups with ADSM.

For detailed information refer to "SQL-BackTrack for Oracle" on page 222 and "SQL-BackTrack for Sybase" on page 227.

SAP/R3

Special interface modules have been specifically developed by the International Technical Support Organization (ITSO) in Boeblingen, Germany for the SAP/R3 environment for Oracle (BACKINT/ADSM) and Adabas (ADINT/ADSM) databases to provide:

- Seamless integration between database administration and ADSM

- Maximum data rates for backup and restore

- Production function necessary in the R/3 environment

Backup and restore of the database are performed by the specific utilities under program (backup) or system administrator (restore) control. R/3 files other than database files, such as executables, control files, and end user data, are backed

up directly through the ADSM client. Thus ADSM can perform a complete
backup of the R/3 system, except for the boot image.

For detailed information refer to "ADSM and SAP R/3" on page 237.

ADSM Notes Agent

Lotus Notes is groupware for sharing information. The information for Lotus
Notes is kept in a repository called the *Lotus Notes database* or *NSF database.*

ADSM V2 is the first product to perform online incremental backups for Lotus
Notes data! With other products (and with ADSM V1) you have to stop the use of
the NSF database, and you can only back up the entire database.

ADSM support for Lotus Notes is currently available for OS/2 Notes servers and
clients. Support for Windows NT and AIX is a high priority requirement. The
NSF database can be backed up while the database is live or online. Because
incremental backup is supported, ADSM will back up only the elements of the
NSF database that have changed (such as documents, forms, and document
design templates). You can also restore an individual note or document.

The Lotus Notes Backup Agent is available with the ADSM V2 OS/2 client.

For detailed information refer to "ADSM and Lotus Notes" on page 240.

ADSM and Complementary Products

* Oracle Support
    SQL-BackTrack from DataTools

* Oracle Parallel Edition, Sybase Support
    SQL-BackTrack from DataTools

* SAP R/3 Support
    RPQ from IBM

* NetWare, Windows NT HSM Client Support
    NetWare with Avail HSM from Wang (Formerly NetSpace)
    Windows NT with OPEN/store from Wang
    Available from IBM

* STK Silo Device Drivers for AIX Server
    IBM Network Tape Access and Control System for AIX (NetTAPE)
        or
    ADSM for AIX with StorageTek Tape Library Support Software

---

ADSM and Complementary Products ...

* Digital VMS Support
    ABC Offering by SSSI
    Available from IBM

* Cray UNICOS, Fujitsu UXP Support
    Offerings from IBM

* Tandem Guardian Support
    BackHome! Offering from ETI

* Client/Server Archival and Retrieval System for Documents
    IBM OnDemand for AIX Version 2.1
    Program Offering from IBM

ADSM and Complementary Products

This graphic lists some products that use the ADSM API. For detailed information refer to Chapter 11, "ADSM and Complementary Products" on page 171.

ADSM and Tivoli

Tivoli provides world class systems management software. ADSM provides the storage management support for a Tivoli environment. Together, ADSM and Tivoli facilitate enterprise-wide systems management and storage management.
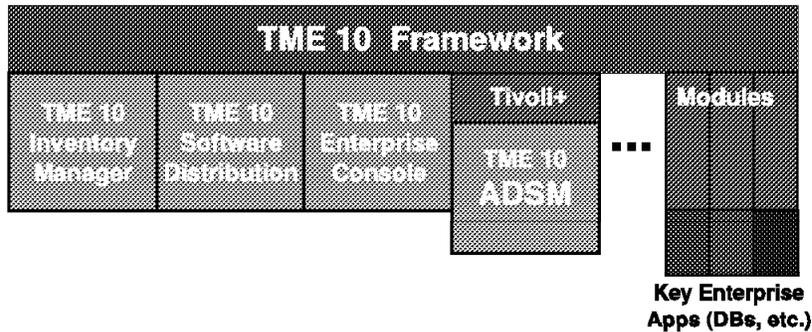
Storage Management a Key TME 10 Application

TME 10 is a new product line from Tivoli that unifies TME 3.0, SystemView, and industry solutions into a true enterprise management solution. Storage management is a key part of TME 10.

Tivoli Provides a Single View, Point of Control into the Enterprise

Tivoli enables you to manage a diverse, distributed enterprise from a single point of control.

Tivoli/Plus Module Connects ADSM to Tivoli

The Tivoli/Plus module integrates ADSM into a Tivoli environment.

Tivoli Can Provide Launch, Event Handling, and Client Code Distribution for ADSM

When the Tivoli/Plus module is used with ADSM, it provides the ability to:

- Launch ADSM client and administrator tasks

- Handle ADSM events such as database backup complete or a media error encountered on the server

• Distribute ADSM client code

For more information about Tivoli, see Chapter 3, "Tivoli/Plus for ADSM" on page 87.

ADSM Information Sources

Online

- **ADSM Homepage:**
  http://www.storage.ibm.com/adsm
- **Network Storage Manager:**
  http://www.storage.ibm.com/nsm
- **Client Code and Updates:**
  index.storsys.ibm.com (primary California)
  ftp.rz.uni-karlsruhe.de (mirror Germany)
  ftp.wu-wien.ac.at   (mirror Austria)
  ftp.cac.psu.edu   (mirror Pennsylvania)
- **Technical Q&A:** IBMLINK, adstarsw@vnet.ibm.com
- **CompuServe - > GO IBMSTORAGE ADSM**
- **ADSM-L Discussion List on the Web**
- **ADSM Consultant**
  http://www.storage.ibm.com/selladsm

FAX or Phone
- **IBM FAX or Phone Services**

Hardcopy Publications
- **Product Publications**
- **ITSO Redbooks**

ADSM Information Sources

This graphic lists the many sources of ADSM information available to you.

Online

- ADSM Homepage

  Starting from this page, you will find detailed information about ADSM, including the latest announcements, news, product information, demos and betas, alliances and partners, and technical support.

- Network Storage Manager

  Information about the new IBM 3466 Network Storage Manager. You can run a demo and view the online information sources.

- Client Code and Updates

  You can find the latest client code updates on the primary FTP server in California for download. Three additional FTP servers are mirror copies of the primary. Feel free to use them.

- Technical Q&A

  You can participate on the IBMLINK ADSMforum if you subscribe (call 1-800-547-1283). On the main menu select Special Interest Area and then ADSMforum.

  You can send presales questions to the ADSM technical support group, using this e-mail address: adstarsw@vnet.ibm.com.

- CompuServe

If you are on CompuServe, use GO IBMSTORAGE ADSM. To subscribe to CompuServe, call 1-800-543-3388.

- ADSM-L Discussion List on the Web

  To participate in this ADSM forum, send e-mail to listserv@vm.marist.edu. In the body of the note say ″SUBSCRIBE ADSM-L yourfirstname yourlastname.″

- ADSM Consultant

  ADSM Consultant is a tool that provides cost justification information and reports for ADSM.

FAX or Phone

- For IBM FAX information services, call 1-800-IBM-4FAX (inside the U.S. and Canada) or 1-415-855-4329 (outside the U.S.).
- For IBM phone information services, call 1-800-IBM-3333 and ask for ″STAR 20.″

Hardcopy Publications

ADSM hardcopy publications are provided with the product. See "ADSM Product Publications" on page 334 for a list of ADSM publications.

ADSM Redbooks

General Topics:
- ADSM Concepts   SG24-4877
- ADSM Version 2 Presentation Guide   SG24-4532
- ADSM Implementation Examples   GG24-4034
- ADSM Advanced Implementation Experiences   GG24-4221
- Using ADSM HSM   SG24-4631

Specific Server Books:
- Getting Started with ADSM/6000   GG24-4421
- ADSM for AIX: Advanced Topics   SG24-4601
- AIX Tape Management   SG24-4705
- ADSM/6000 on 9076 SP2   GG24-4499
- ADSM for MVS: Recovery and Disaster Recovery   SG24-4537
- ADSM for MVS: Using Tapes and Tape Libraries   SG24-4538
- Getting Started with ADSM/2   GG24-4321
- ADSM for OS/2: Advanced Topics   SG24-4740
- Setting Up and Implementing ADSM/400   GG24-4460
- ADSM/VSE Implementation Guide   SG24-4266

Specific Client Books:
- Getting Started with ADSM NetWare Clients   GG24-4242
- Getting Started with ADSM AIX Clients   GG24-4243
- ADSM API Examples for OS/2 Windows   SG24-2588

ADSM with Other Products:
- Using ADSM to Back Up Databases   GG24-4335
- Using ADSM to Back Up Lotus Notes   SG24-4534
- HSM for NetWare: ADSM and Available HSM Implementation   SG24-4713
- Using ADSM to Back Up OS/2 LAN Server and WARP Server   SG24-4682
- Backup, Recovery, and Availability with DB2 Parallel Edition   SG24-4695

Coming Soon:
- ADSM Server for NT Configuration and Recovery Examples   SG24-4878
- Client Disaster Recovery: Bare Metal Restore   SG24-4880

ADSM Redbooks

The graphic lists all IBM Redbooks related to ADSM. For up-to-date information about all redbooks on ADSM and other IBM products, refer to the redbooks homepage:

http://www.redbooks.ibm.com/redbooks

or

http://www.redbooks.ibm.com/redbooks/findrbk.html

Implementing ADSM

IBM provides ADSM services that are complete solutions, or you can select components you need. There are also IBM business partners that can help you. ADSM implementation will vary according to customer needs and home location (country). If you need more specific information about ADSM implementation offerings, please contact your local IBM marketing and sales representative.

IBM 3466 Network Storage Manager

Network Storage Manager integrates leading IBM software and hardware storage technologies into a single powerful solution.
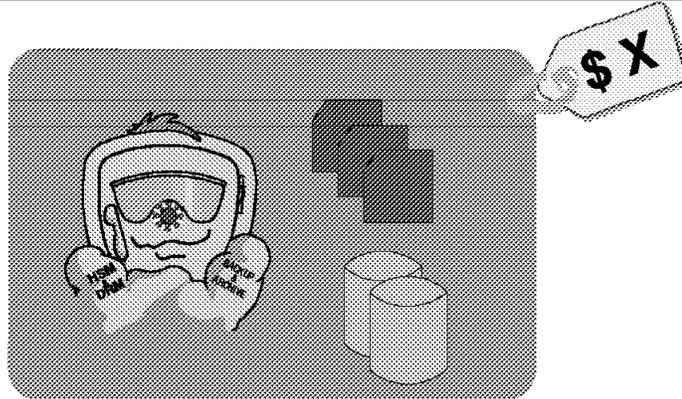
Software- Hardware Integrated

You not have to spend your valuable time evaluating, integrating, and testing separate products to devise a solution, so you get more for your investment, and you have a single point of contact for all of your service and support needs. Backed by IBM know-how and quality, Network Storage Manager is designed to help you protect your most valuable strategic asset—your business data.

Network Enabled

The Network Storage Manager is configured, loaded, and ready to plug in to your network. It comes with attachments for Token Ring, Ethernet, and Single-Ring and Double-Ring FDDI.

Basic Customization Included

All delivered components have been preconfigured and pretested. Almost everything a basic user requires is provided. If you choose to use a special communications protocol or want customization, additional work will be required. You can either perform the customization yourself or arrange for IBM or an IBM Business Partner to assist you on a fee basis.

″Dream Team″ of First-Rate Components

IBM combines these industry-leading hardware and software components into a tightly integrated offering:

- Award-winning ADSTAR Distributed Storage Manager for AIX V2.1 software

- High-speed Serial Storage Architecture (SSA) disk storing 72 to 144 GB of data

- Reliable 1/2-inch Magstar Tape Library holding 18.3 to 40.2 TB of compressed data

- Powerful RISC processor

One-Price, One-Stop Shopping

The end result? A powerful, one-price one-stop solution that greatly simplifies the entire range of storage management tasks for end users as well as storage and network administrators, no matter where they work or on which systems they rely to get their jobs done.

Different Configurations Available

Different configurations are available to suit customer needs.

For more information about the Network Storage Manager, refer to Chapter 7, "Network Storage Manager" on page 125.

Summary
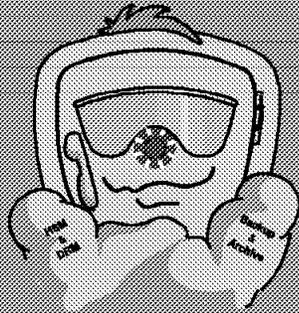
Cover Your Assets with ADSM!

ADSM is a scalable, flexible, enterprisewide, perfect solution for heterogeneous environments.

Remember what ADSM can do for you:

- Supports many clients, servers, and databases
- Works with many storage devices
- Provides performance and function
- Portable between platforms
- Reasonably priced

Chapter 2. Disaster Recovery Manager

Disaster Recovery Manager Overview
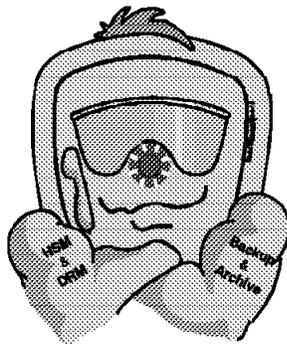
The Disaster Recovery Manager (DRM) is an optional ADSM feature that facilitates an ADSM-based recovery of business applications from backup data stored offsite. Recovery can potentially be performed at an alternative site, on replacement hardware, by people not familiar with the backup applications.

The DRM disaster recovery plan is used to guide you through disaster recovery as well as provide an auditable procedure to certify the recoverability of the ADSM server. DRM provides automated generation of the server disaster recovery plan file, offsite recovery media management, and storage of client recovery information.

In this chapter we cover the topics listed on this graphic.

Availability Management and
Disaster Recovery Management
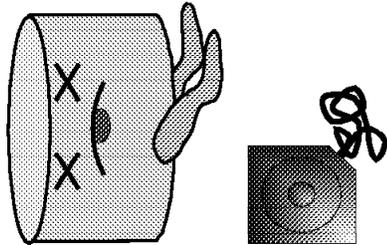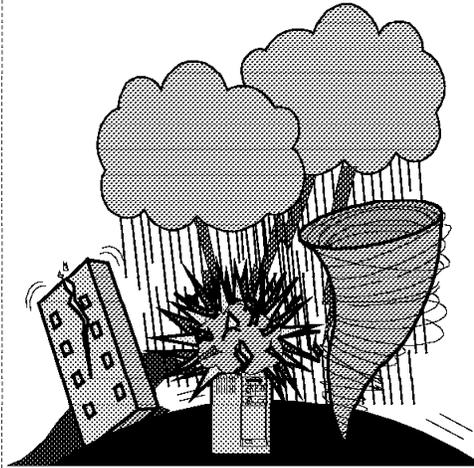
* Availability
Management

* Disaster
Recovery
Management

Availability Management and Disaster Recovery Management

Let's compare availability management with disaster recovery management.

Availability Management

Availability management is used for recovery from incidental computer system outages such as disk drive or tape crashes. Downtime is often minimized by using disk mirroring and other forms of RAID technology or by maintaining onsite backup copies of data.

Availability management for the ADSM server is accomplished with ADSM by:

- Mirroring the server database and recovery log

- Backing up storage pools and storing them onsite

Disaster Recovery Management

A disaster is a catastrophic interruption of business processing that destroys the ADSM server, or ADSM clients, or both. Backup data is located offsite to protect it from damage.

Disaster recovery management is accomplished with ADSM by:

- Backing up client data to the ADSM server

- Saving the client recovery information in the ADSM database

- Backing up the primary storage pools and storing the media offsite

- Backing up the server database to removable media and storing the media offsite

- Generating the disaster recovery plan file to assist with ADSM server recovery

Backing up the storage pools will introduce changes to the database, so it is important to back up the primary storage pools before backing up the database.

Benefits

Reduces Cost and Simplifies
Disaster Recovery Planning

Automates Maintenance of
Recovery Plan

Minimizes Downtime and
Errors

Disaster
Recovery
Manager

Benefits

You certainly know that disaster recovery is a very complex procedure, and you are increasingly asked to develop more recovery plans with less effort and fewer resources. DRM facilitates disaster recovery in three ways:

Reduces Cost and Simplifies Disaster Recovery Planning

The cost of implementing disaster recovery is one of the factors to consider when you decide whether or not to implement DRM. You also must consider how much it will cost you not to implement DRM.

DRM is the tool you need to recover from a disaster. It provides a well-documented plan with the steps required for a successful recovery. You do not have to worry about which steps to execute or in which order. DRM makes those decisions for you, thus saving you time as well as money.

Automates Maintenance of Recovery Plan

Because the preparation of the disaster recovery plan can be automated to run at appropriate intervals, your disaster recovery plan is always up-to-date. During the preparation step, all required information is collected and stored in one place, the disaster recovery plan file.

The disaster recovery plan file provides an auditable procedure for ADSM server recovery.

Minimizes Downtime and Errors

From the disaster recovery plan you get a complete list of the backup volumes you have to move from the offsite to the onsite location. DRM helps you minimize the time it takes to obtain the correct backup volumes, thereby preventing errors during the recovery process.

Remember, the better the recovery plan, the easier the recovery!
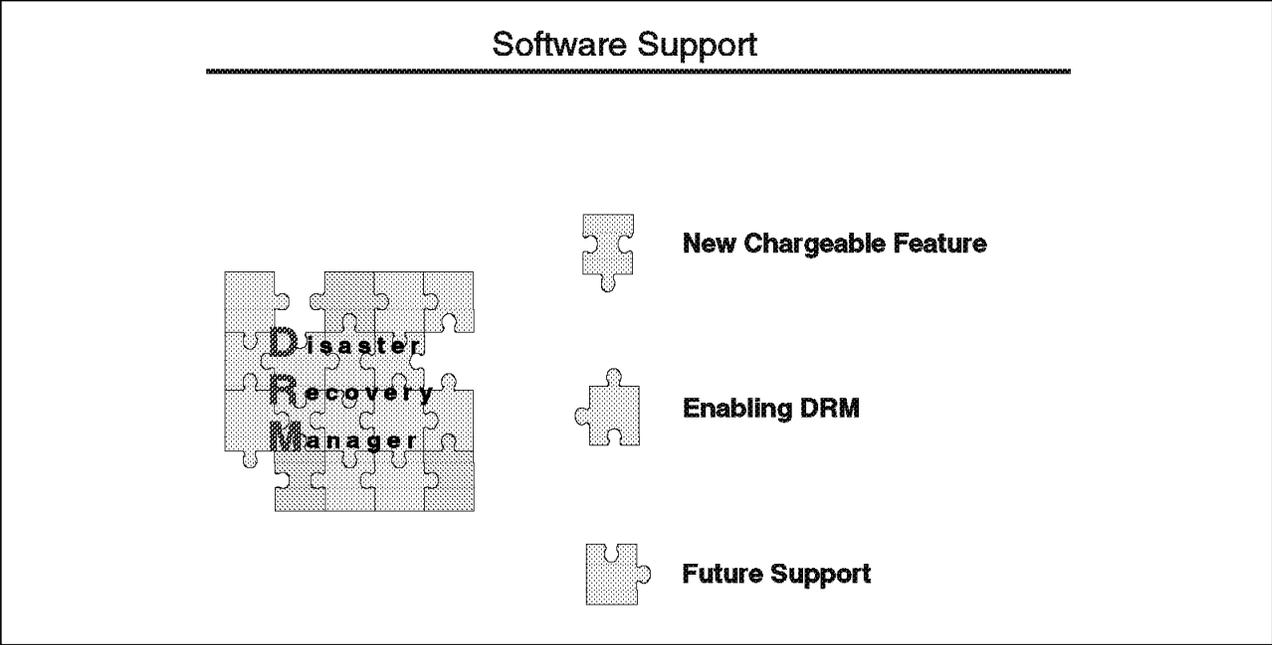
Software Support

Software Support

DRM is available for the ADSM server on the MVS and AIX platform.

New Chargeable Feature

DRM is a new chargeable feature of ADSM for AIX Version 2 and ADSM for MVS Version 2.

The code for DRM is shipped with the ADSM server. A new ADSM customer gets the code with the ADSM server. A customer that already has ADSM installed gets the code through an ADSM server PTF.

Enabling DRM

Use the REGISTER LICENSE command to enable the ADSM server to support the DRM feature.

Future Support

IBM intends to provide DRM on other ADSM server platforms.
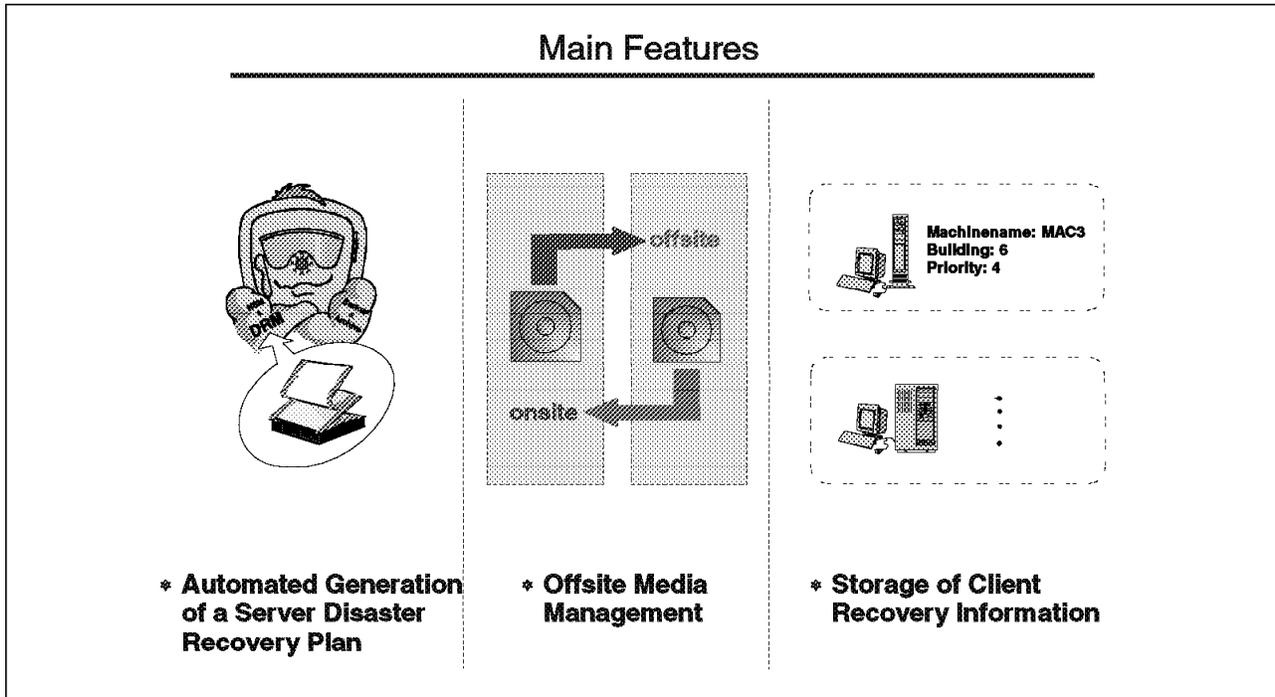
Main Features

Main Features

The main features of DRM are:

- Automated generation of a server disaster recovery plan

- Offsite media management

- Storage of client recovery information

Automated Generation of a Server Disaster Recovery Plan

The PREPARE command automatically queries the required information from the ADSM server and generates a recovery plan file that is based on a predefined recovery strategy for the server. The PREPARE command is scheduled through ADSM central scheduling to maintain an up-to-date recovery plan.

The disaster recovery plan file contains the information and procedures necessary to assist you with the recovery of the ADSM server.

For more information see "Disaster Recovery Plan File" on page 80.

Offsite Media Management

Knowing the location of offsite recovery media is critical to successful disaster recovery. With ADSM you can determine which database backup and copy storage pool volumes have to be moved from the offsite location and back onsite. ADSM keeps track of the location of volumes.

Database backup and copy storage pool volumes are treated as logical collections that are selected to be moved offsite for safekeeping and onsite for reuse or disposal. The reclamation of offsite volumes enables you to perform expiration of an ADSM database backup series.

DRM offsite recovery media management can work in conjunction with or independently of a tape management product.

For more information see "Offsite Media Management" on page 82.

Storage of Client Recovery Information

You use ADSM commands to store information about a machine where an ADSM client resides. DRM allows the following client recovery information to be saved in the ADSM database:

- Business priority
- Machine locations, machine characteristics, and machine recovery instructions
- Boot media requirements

In the event of a disaster, DRM QUERY commands provide assistance to help you determine:

- Which client machines were lost in the disaster and have to be recovered
- The priority of the client machines to identify the recovery order
- The machine requirements and boot media requirements

Disaster Recovery Plan File

* Command Stanzas
* Site-Specific Instruction Stanzas
* Server Requirements Stanzas
* Configuration File Stanzas
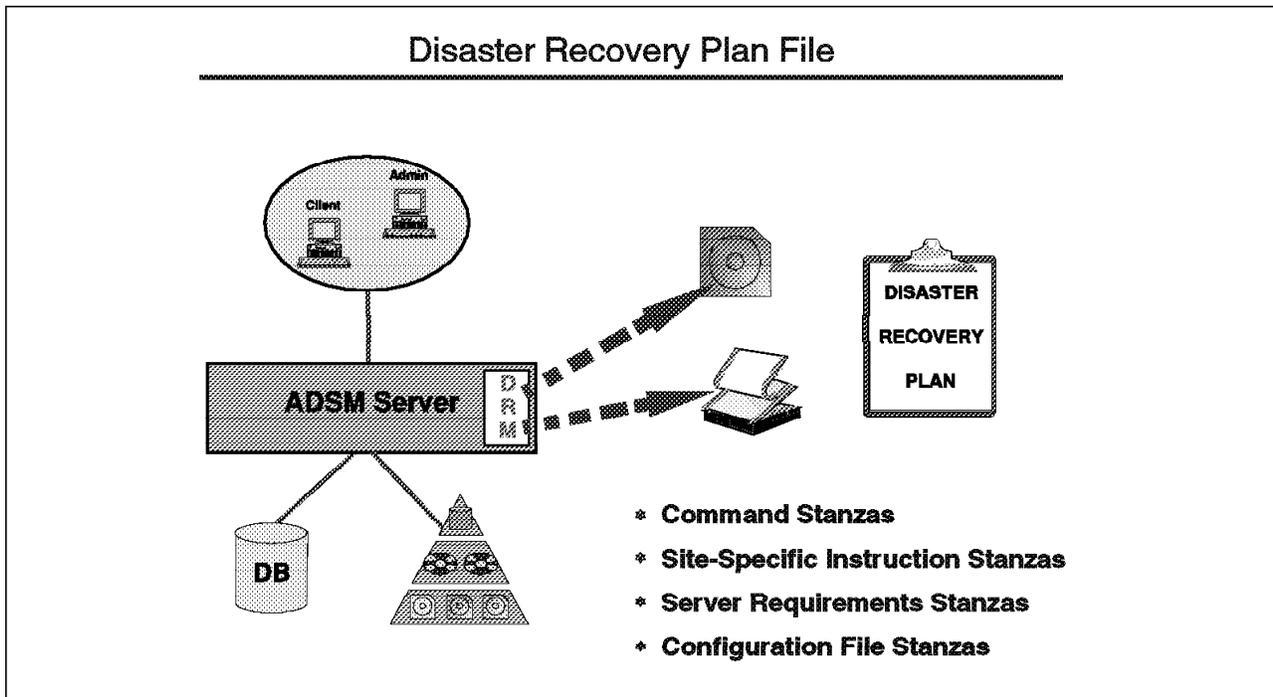
Disaster Recovery Plan File

The DRM disaster recovery plan provides information such as the sequence of steps to follow for the recovery of the ADSM server.

When you issue the PREPARE command, DRM queries the ADSM server for required information to generate a disaster recovery plan file. The disaster recovery plan file includes all required recovery information in one place.

To ensure that the disaster recovery plan file is kept up-to-date, run the PREPARE command after the database backup and storage pool volumes have been marked to be sent offsite.

The recovery information is organized into stanzas in the disaster recovery plan file. These stanzas can be divided into four categories:

- Command stanzas
- Site-specific instruction stanzas
- Server requirements stanzas
- Configuration file stanzas

Command Stanzas

For the MVS platform the command stanzas provide the JCL, REXX EXECs, and ADSM macros. For the AIX platform, the command stanzas provide the shell scripts and ADSM macros. You can view, print, update, or execute the command stanzas as part of the recovery process.

The command stanzas provide commands to:

- Create and format the database and recovery log volumes

- Restore the database

- Start the server

- Tell the server which volumes are available for use

- Tell the server which primary storage pools have been destroyed

- Create and format the new storage pool volumes

- Restore the primary storage pools from the copy storage pools

Site-Specific Instruction Stanzas

Site-specific instruction stanzas consist of recovery instructions specific to your site. You can print, update, and use these stanzas during server recovery.

Site-specific instruction stanzas contain the following information:

- Name of the person to contact for the recovery process

- ADSM administrator ID and password

- Vault location where the offsite volumes reside

- Courier's name and telephone number

- ADSM installation procedures, name of the install volume, and license numbers

- How to recover the database and find disk space for the database and recovery log

- Order in which the storage pools should be recovered

Server Requirements Stanzas

Server requirements stanzas consist of the database and recovery log requirements and volume and device requirements. You can view or print these stanzas.

Server requirements stanzas contain the following information:

- Space requirements for the database and recovery log

- Volumes required for restore of the database and storage pools

- Device class definition of the device required to read the volumes that are used to restore the database

Configuration File Stanzas

Configuration file stanzas consist of the volume history, device configuration, and server options files.

For a more detailed description of the stanzas, see the *ADSM for MVS Disaster Recovery Manager Administrator's Guide and Reference* (SG35-0238) or the *ADSM for AIX: Administrator's Guide* (SH35-0134).

Offsite Media Management
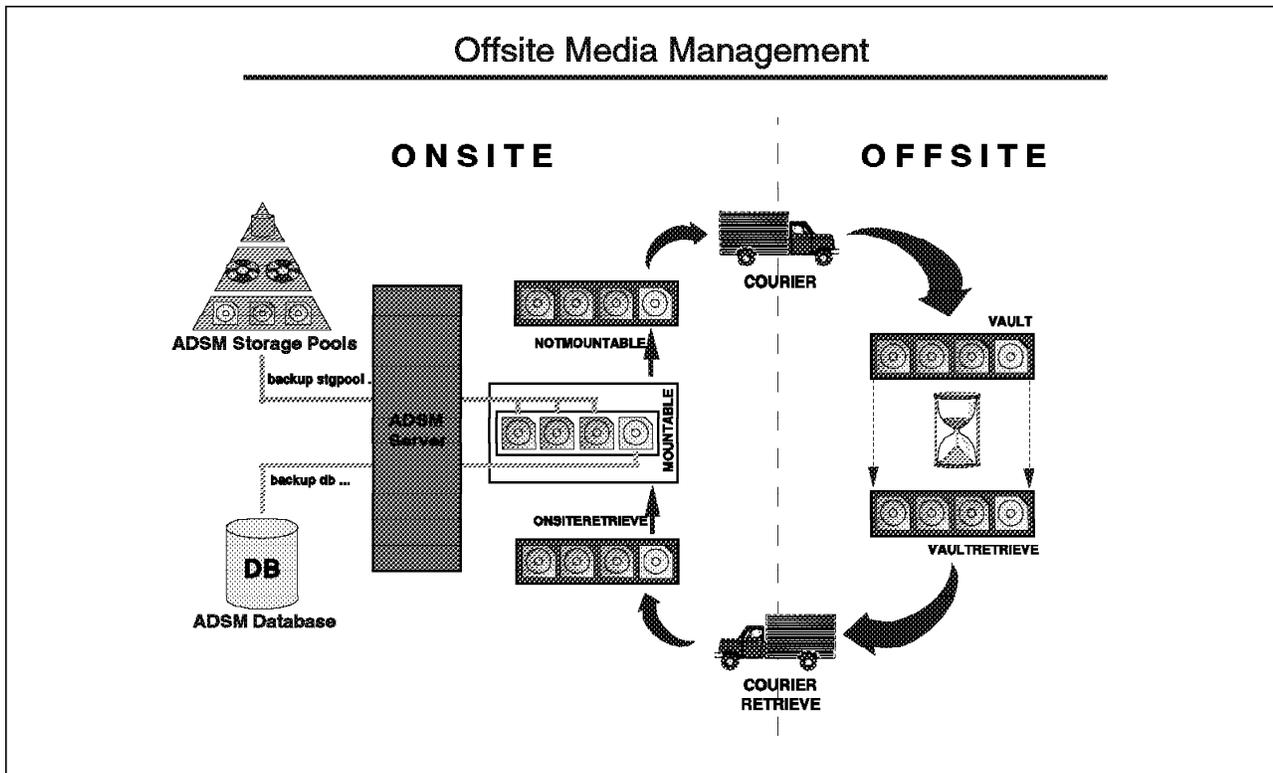
Offsite Media Management

Offsite media management is used during routine operations and defines a process for moving ADSM database backup and copy storage pool volumes:

• Offsite for disaster recovery protection

• Onsite when they no longer contain valid data

Issue the MOVE DRMEDIA command to indicate the movement of the volumes and the QUERY DRMEDIA command to display and keep track of the location of the volumes.

The PREPARE command generates the backup volume information in the disaster recovery plan file. If the ADSM server is destroyed, the disaster recovery plan provides you with a list of offsite volumes required at the recovery site.

Sending Server Backup Volumes Offsite

DRM uses four states for database backup and copy storage pool volumes that are sent offsite for disaster recovery protection:

• Mountable - The volume contains valid data, is accessible for the ADSM server.

• Notmountable - The volume contains valid data and is unavailable to the ADSM server, but is still onsite.

• Courier - The volume contains valid data and is with the courier or shipping company.

- Vault - The volume contains valid data and is at the vault.

You can use these states in any order. For example, you can use the MOVE DRMEDIA command to indicate a movement from the Notmountable state to the Vault state, or you can indicate a movement from the Courier state to the Mountable state.

Moving Reclaimed or Expired Volumes Back Onsite

DRM uses three states for moving back on site backup volumes that have been reclaimed or no longer contain valid data:

- Vaultretrieve - The volume no longer contains valid data and is to be returned. The vault operator should give the volume to the courier.

- Courierretreive - The volume no longer contains valid data and is in the process of being returned by the courier.

- Onsiteretrieve - The volume no longer contains valid data and has been returned to the onsite location. The volume records of ADSM database backup and scratch copy storage pool volumes are deleted from the ADSM database. For private copy storage pool volumes, the access mode is updated to READWRITE.

For a more detailed description of offsite media management, see the *ADSM for MVS Disaster Recovery Manager Administrator's Guide and Reference* (SG35-0238) or the *ADSM for AIX: Administrator's Guide* (SH35-0134).
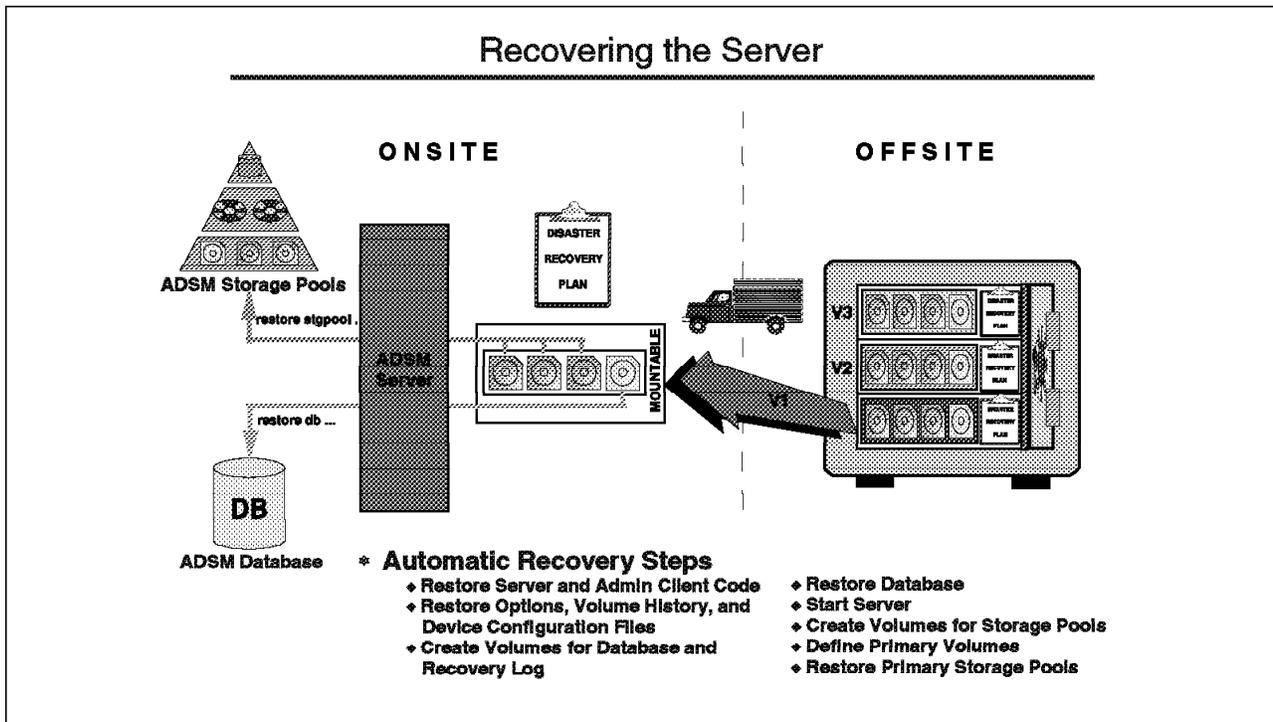
Recovering the Server

After a disaster the ADSM server requires a complete recovery. To recover the ADSM server, you need the disaster recovery plan file, the backup volumes, and a replacement machine for the ADSM server.

Obtain the latest recovery plan file and, if necessary, update the file.

Obtain the backup volumes for the database and storage pools from the offsite location.

Locate a suitable replacement machine for the ADSM server. Information such as memory size and machine type as well as install instructions for the ADSM server can be stored in one of the stanzas in the disaster recovery plan file.

Automatic Recovery Steps

Use the disaster recovery plan to recover the ADSM server. The stanzas in the disaster recovery plan file automatically perform the steps listed below. Only minimal editing is required.

1. Restore the ADSM server and the administrative client software.

2. Restore the ADSM server options, volume history, and device configuration files.

3. Create the ADSM server volumes for the database and recovery log.

4. Restore the ADSM database.

5. Start the server.

6. Create the ADSM server volumes for the new primary storage pools.

7. Define the primary volumes to the ADSM server.

8. Restore the primary storage pools from the backup volumes.

Test of the Recovery Procedures

It is important to test your recovery procedures. We strongly recommend that you periodically perform a recovery test of your ADSM server.

For a more detailed description of recovering the server, see the *ADSM for MVS Disaster Recovery Manager Administrator's Guide and Reference* (SG35-0238) or the *ADSM for AIX: Administrator's Guide* (SH35-0134).
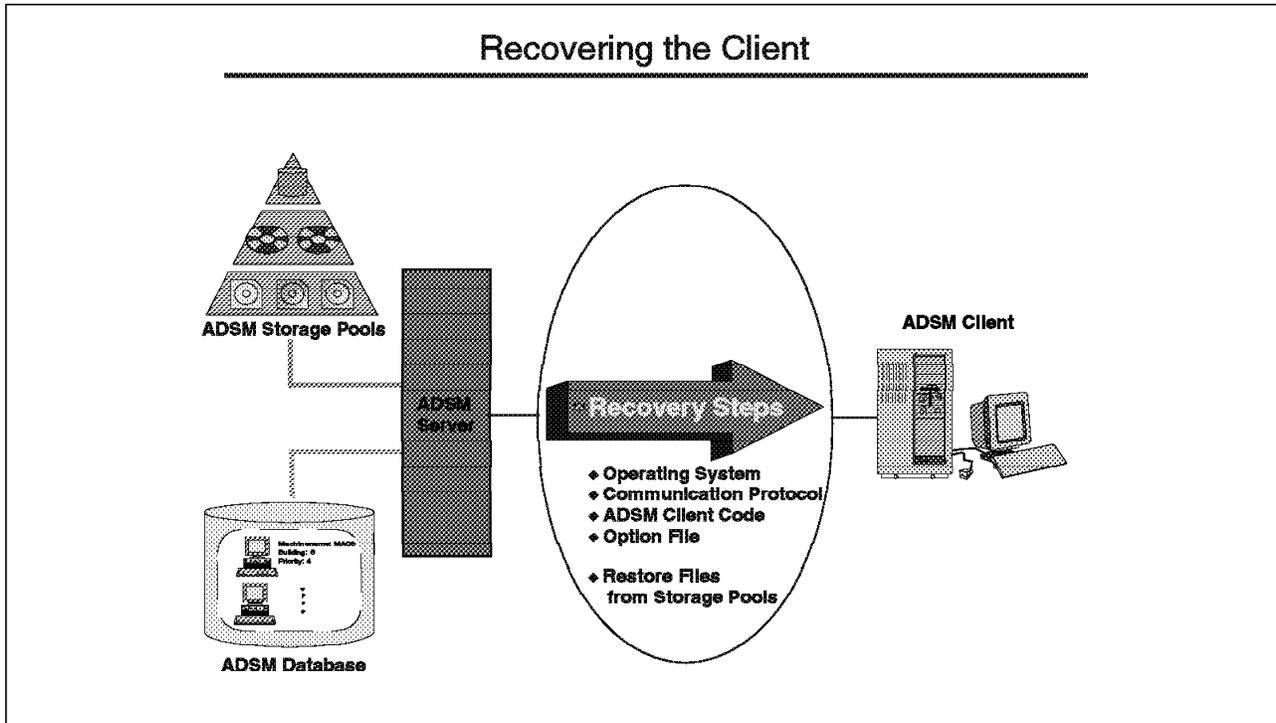
Recovering the Client

Using DRM, you can store information about the ADSM clients, such as
hardware and software requirements, in the ADSM database. When the
information about the ADSM client is defined and stored, it can be a part of the
disaster recovery procedure, and the available information can be used to
restore an ADSM client. First you have to:

- Determine which client nodes were lost in the disaster

- Identify in which order the clients should be recovered

- Find a replacement machine and the boot media

Recovery Steps

Once you have found a suitable replacement machine for the ADSM client, use
the QUERY command to obtain information about the machine. After you install
the hardware, follow these steps:

1. Restore the operating system software.

2. Restore the communication protocol.

3. Restore the ADSM client code.

4. Restore the ADSM client options files.

5. Restore the client file systems from the primary storage pools.

For more information about recovering an ADSM client, see *ADSM Client
Disaster Recovery: Bare Metal Restore* (SG96-0402) and the *ADSM for MVS
Disaster Recovery Manager Administrator′s Guide and Reference* (SG35-0238) or
the *ADSM for AIX: Administrator′s Guide (SH35-0135).*

Chapter 3. Tivoli/Plus for ADSM

TME 10 Overview

TME 3.0

Tivoli/Enterprise Console
Tivoli Management Framework

SystemView

NetView
System Monitor/6000
OPC
ADSM
DCAF
...

Other Products

Maestro
BoKS
...

Tivoli/Enterprise Console	Tivoli Management Framework	TME 10 Base Products	Tivoli/Plus Modules	3rd Party Applications

TME 10

TME 10 Software Distribution   Tivoli/Plus for ADSM
TME 10 Distributed Monitoring   Tivoli/Plus for BoKs
TME 10 Job Scheduling   ...
TME 10 NetView
TME 10 Performance Management
TME 10 Print Management
TME 10 ADSM
...

TME Overview

TME 10 is a new product line from Tivoli that brings together TME 3.0, SystemView, and industry solutions into a true enterprise management solution. Tivoli now unifies TME and SystemView products around a distributed management technology, the Tivoli Management Framework (TMF). The industry solutions consist of other systems management solutions, hardware, software, databases, or applications that are compatible with TME 10.

The consolidation of the TME 10 product line is not a process of choosing one offering over another but of providing a real distributed systems management solution.

TME 3.0

TME 3.0 is a comprehensive set of management applications, integrated into a distributed object framework. It provides control and monitoring for the managed resources through a common GUI desktop and integrated event console (Tivoli/Enterprise Console or T/EC).

T/EC offers a centralized view of the distributed computing enterprise but also ensures high availability of applications and computing resources. T/EC collects, processes, and automatically responds to common management events, such as a database server being down, a network connection being lost, or a batch processing job having completed successfully. It acts as a central collection point for alarms and events from all key sources.

The TMF is an object-oriented architecture and constitutes the core of TME. The TMF includes services for transactions, security, events and alerts, configuration, and automation based on platform-independence. So it can manage any object, regardless of type (UNIX, OS/2, NT, database, or application). The TMF provides all necessary managers, brokers, and agents to ensure seamless support for heterogeneous servers and clients, coordinated operations, replicated services and servers, high flexibility, and an open set of APIs and services based on key standards.

SystemView

Compared with Tivoli, SystemView brings a group of management functions matching each of IBM's operating platforms: MVS, OS/400, AIX, and OS/2. Each platform has specific functions and specific interfaces, such as RACF for MVS and Job Scheduler for AIX. TME 10 currently integrates SystemView for AIX and SystemView for OS/2.

Tivoli has sought partners for enterprise management tasks such as performance management, help desk, backup/restore, and network management. These partnership activities take place under the auspices of the Tivoli/Plus program. Some of these products were already part of the SystemView platform, such as ADSM, but others like BoKS and Maestro are non-IBM products.

TME 10

TME 10 consists of several products. The goal of this presentation is not to describe all of them, only those that are important for a Tivoli/Plus for ADSM installation.

Tivoli/Enterprise Console

T/EC is a management-automation application. It collects information from all sources (Tivoli applications, Tivoli partners, network management platforms, relational database systems, customer applications, and heterogeneous operating systems). T/EC's grouping and filtering capabilities reduce the number of events displayed to the operations staff.

T/EC event adapters capture and filter messages and then forward them as events to the T/EC console.

Tivoli Management Framework

The TMF supports scalable, reliable, and secure distributed management applications. It models every resource in the distributed enterprise (clients, servers, database applications, and users) as objects. These objects provide a level of abstraction that hides the complexity of managing large numbers of heterogeneous resources.

The TMF is an open, object-oriented architecture that constitutes the core of TME 10 management by including the necessary set of managers, brokers, and agents. The key services that the TMF provides are shared by all Tivoli applications.

TME 10 Base Products

The following are some examples of TME 10 base products:

TME 10 Software Distribution

TME 10 Software Distribution provides a framework of services and applications for the distribution and maintenance of operating systems, applications, and data across networks of workstations that run on a variety of software platforms. TME 10 Software Distribution provides extensive facilities for deploying software to distributed systems. Currently it offers you the choice of implementing Tivoli/Courier or IBM Software Distribution for AIX. Because these products support most UNIX and PC operating systems as clients, they can be helpful when deploying ADSM code throughout the network.

TME 10 Distributed Monitoring

TME 10 Distributed Monitoring provides an automated way of maximizing the availability of key, business-critical computing resources in distributed environments. Designed specifically for client/server computing, TME 10 Distributed Monitoring consists of the Tivoli/Sentry and Tivoli System Monitor products.

TME 10 Distributed Monitoring:

- Groups and monitors distributed computing resources
- Automatically detects and corrects problems
- Provides notification and escalation of problems when necessary
- Provides central, consistent control for managing unattended remote sites

TME 10 Distributed Monitoring provides central configuration of monitoring parameters (prepackaged or user defined) at remote sites from a central location and offers practical, ready-to-use monitors. It enables centralized information technology staff to set or change monitoring parameters for hundreds of related, remote systems with a single operation. Tivoli/Sentry provides proactive management of client/server resources by automatically initiating predefined preventive and corrective actions.

The installation of TME 10 Distributed Monitoring is useful for automating responses to ADSM messages.

TME 10 ADSM

TME 10 ADSM is the ADSM product itself, when ordered through Tivoli. You do not have to order TME 10 ADSM specifically if you want to integrate ADSM into a Tivoli platform. ADSM for AIX can be integrated into Tivoli too. In both cases the mandatory product for such an integration is the Tivoli/Plus for ADSM module.

Tivoli/Plus Modules

Tivoli/Plus is a powerful collection of ready-to-use software modules that link a product application (for example, ADSM) with TME 10. Many modules are already available and more are on the way.

Third Party Applications

In some cases, third party vendors have decided to write object-oriented applications directly on the TMF to port them to the many operating systems on which the TMF runs. This integration enables function sharing, such as the use of a common TMF scheduler or TMF policy manager. Examples of third party applications that exploit the TMF are those that manage Oracle and Sybase databases. These applications are not included in the TME 10 products

Other TME 10 Products and Tivoli/Plus Modules

Please refer to announcement letter ZP96-0411 (296-216 in the U.S.) for more information about the following TME 10 products and Tivoli/Plus modules:

- TME 10 Event Integration Facility
- TME 10 NetView
- TME 10 Performance Management
- TME 10 Performance Reporter
- TME 10 User Administration
- TME 10 Remote Control
- TME 10 Print Management
- Tivoli/Plus for AR System
- Tivoli/Plus for ServiceCenter
- Tivoli/Plus for BoKS
- Tivoli/Plus for SeOS
- Tivoli/Plus for Maestro
- TME 10 Application Development Environment
- TME 10 Application Extension Facility
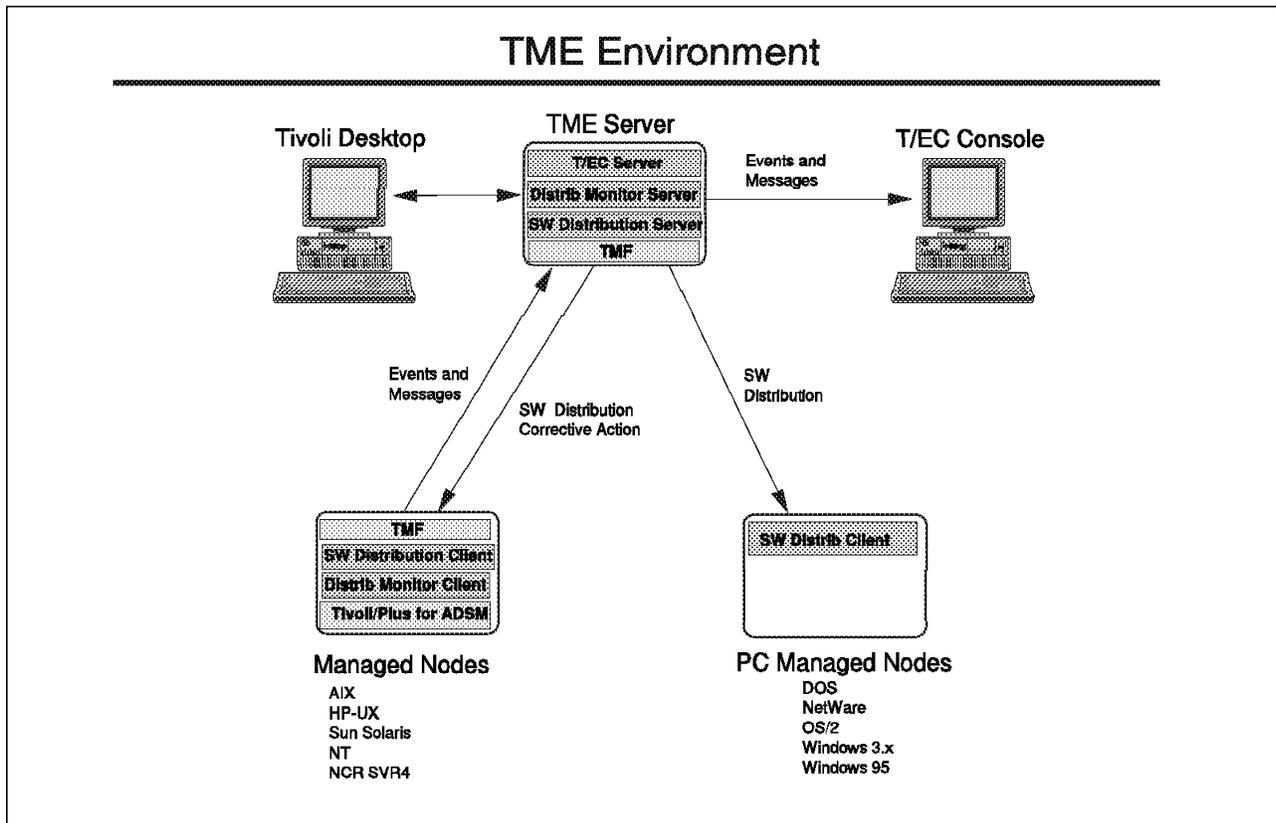
TME Environment

TME Environment

This graphic is a simplified representation of a TME environment. Each of these component is described in detail on the pages that follow.

TME Server

The TME server manages the TME environment and keeps policy information. The TME server is also referred to as the Tivoli Management Region (TMR) server.

TMF

The Tivoli Management Framework is supported on UNIX and NT platforms and provides common services to TME servers and managed nodes.

T/EC Server

The T/EC server receives events sent from managed nodes. The T/EC server records the messages and events in a database and sends them to the T/EC console. Corrective actions can be automatically performed according to rules that have been established on the TME server.

This graphic shows the T/EC server on the same machine as the TME server. The T/EC server can also run on a managed node.

Distributed Monitor Server

The distributed monitor server receives monitored events and sends them to the Tivoli Desktop and the T/EC console.

Software Distribution Server

The software distribution server is responsible for deploying code to managed nodes and PC-managed nodes.

Tivoli Desktop

The **Tivoli desktop** allows a Tivoli administrator to perform systems management tasks. The Tivoli desktop is used to access Tivoli applications and Tivoli/Plus modules and can run on any managed node or Windows 95 system.

T/EC Console

The T/EC console receives messages from the T/EC server and can run on any managed node or Windows 95 system. T/EC can be tailored so that different T/EC consoles receive different types of messages. The Tivoli Desktop is used to start the T/EC console. Thus the Tivoli Desktop must be on the same machine as the T/EC console.

Managed Nodes

Managed nodes are machines on which the TMF is installed. TME applications such as TME 10 Distributed Monitoring (Sentry) use TMF services to operate in a TME environment.

The Tivoli/Plus for ADSM module is installed on a managed node, which can be on an administrative, backup-archive, or HSM client on a UNIX or NT platform. When managing an ADSM server, the Tivoli/Plus module uses the administrative client interface. Therefore, you can manage any ADSM server as long as you connect to it from a UNIX or NT administrative client.

PC-Managed Nodes

To provide full function and security, most TME applications require the TMF. However, some of them (in particular Software Distribution) can use the limited version of the framework provided by a PC-Managed Node. PC managed nodes support the platforms shown on the graphic.

Tivoli/Plus for ADSM

The Tivoli/Plus for ADSM module provides integration of the Tivoli environment with ADSM. This integration allows centralized distribution and management of ADSM across a multiplatform network.

Tivoli/Plus for ADSM: Software Distribution
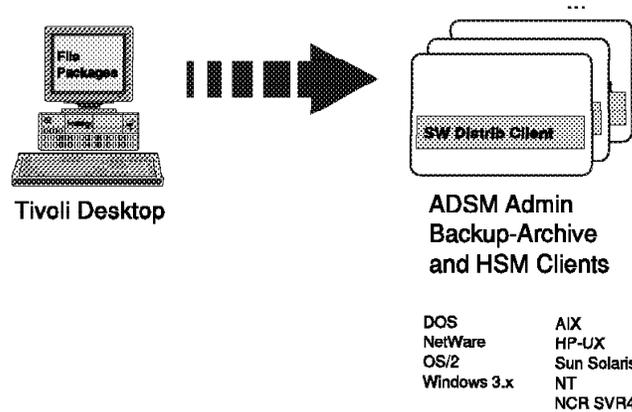
Tivoli/Plus for ADSM: Software Distribution

The Tivoli/Plus for ADSM module enables you to distribute ADSM administrative, backup-archive, and HSM client code. Currently, the module distributes the file packages but does not install them. TME 10 Software Distribution provides a pre- and post- distribution exit that you can use to automatically install the software.

Tivoli Administrator

When the Tivoli/Plus for ADSM module is installed, a Tivoli administrator can use the Tivoli desktop to see the icons for Tivoli/Plus for ADSM. The Tivoli administrator can select ADSM file packages to distribute the ADSM client code across the network.

ADSM Administrative, Backup-Archive, and HSM Clients

You can distribute all components of ADSM client code: administrative, backup-archive, and HSM. Although most TME applications require the TMF, Software Distribution does not. Therefore, UNIX, NT, and non-UNIX clients are supported.

Tivoli/Plus for ADSM: Event Management
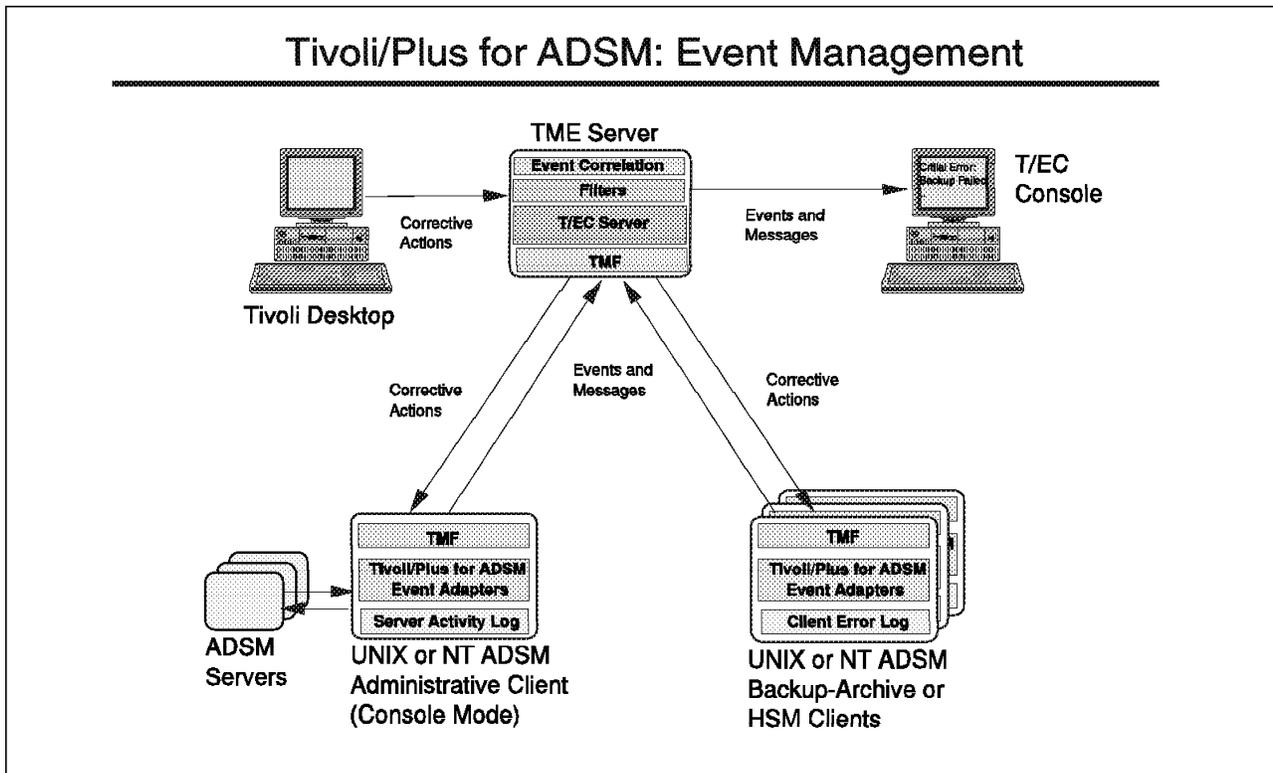
Tivoli/Plus for ADSM: Event Management

The Tivoli/Plus for ADSM module comes with event adapters that filter through messages in the client and server logs for predefined ADSM events.

ADSM UNIX or NT Administrative Client (Console Mode)

Remember that the Tivoli/Plus for ADSM module is installed on a managed node; that is, an administrative, backup-archive, or HSM client on a UNIX or NT platform. Therefore, you can manage any ADSM server as long as you connect to it from a UNIX or NT administrative client.

ADSM UNIX or NT B/A or HSM Clients

You can only perform event management on UNIX or NT backup-archive or HSM clients.

Tivoli/Plus for ADSM Event Adapters

ADSM events are assigned severity levels and sent to the T/EC server. Once event messages are sent to the T/EC server, either the Tivoli administrator can respond appropriately or the messages can trigger automated corrective actions that do not require human intervention. The severity levels help put incoming enterprise events in a manageable format. For example, the system administrator can choose to view only fatal or critical events.

Filters

The T/EC server groups events into classes, which can then be sent to different T/EC consoles. For example, a message about an incorrect volume mount might go to one console, and a message about running out of recovery log space might go to another console.

Event Correlation

T/EC provides a rule engine that can examine the event stream in order to suppress, highlight, or correlate events. Some events may seem harmless if they occur only once, but they can be disastrous if they occur several times a day. Event correlation rules can recognize when recurring events indicate an underlying chronic condition. For example, if someone attempts to log on with an incorrect password once, you need not worry. If, however, someone tries to log on with an incorrect password more than four times on the same day, you might start to worry.

Another example of a recurring event indicating an underlying chronic condition is incremental database backup triggered because the log full percentage was exceeded more than three times within a 24-hour period.

Just like filtered events, correlated events are forwarded to the T/EC server and can automatically trigger corrective actions such as locking a client node that has excessive logon attempts.

A default set of event filters is provided with the Tivoli/Plus for ADSM module. You can tailor these filters to suit your needs.

Tivoli/Plus for ADSM: Distributed Monitoring

If the system on which an ADSM server is running suddenly crashes, there is no way for that managed node to send a message about an ADSM event to anyone.

Two events are currently monitored for ADSM:

- The host system where an ADSM server is running goes down.

- The percentage of network collisions is more than 3% or 4%. This applies to Ethernet networks only.

  If either of these events occurs, a Severe or Critical message is sent to the TME server, which sends a message to the T/EC console and updates the monitors on the Tivoli desktop. The percentage of network collisions more than 3% is considered Severe. Host system down and the percentage of network collisions more than 4% are considered Critical. The monitors on the Tivoli desktop are represented as thermonitors, that is, thermometers.

This graphic shows the TME server and T/EC server are shown on different machines. They could also run on the same machine.

Tivoli/Plus for ADSM: Tasks and Jobs

Tasks and Jobs

TME Server

Tivoli Desktop

Add DB Volume
Add Log Volume
Register Node
Query ....
...

Archive
Incremental Backup
Restore Backup
Query...
...

TMF
Tivoli/Plus for ADSM

TMF
Tivoli/Plus for ADSM

ADSM
Servers

UNIX or NT ADSM
Administrative Client

UNIX or NT ADSM
Backup-Archive or
HSM Clients
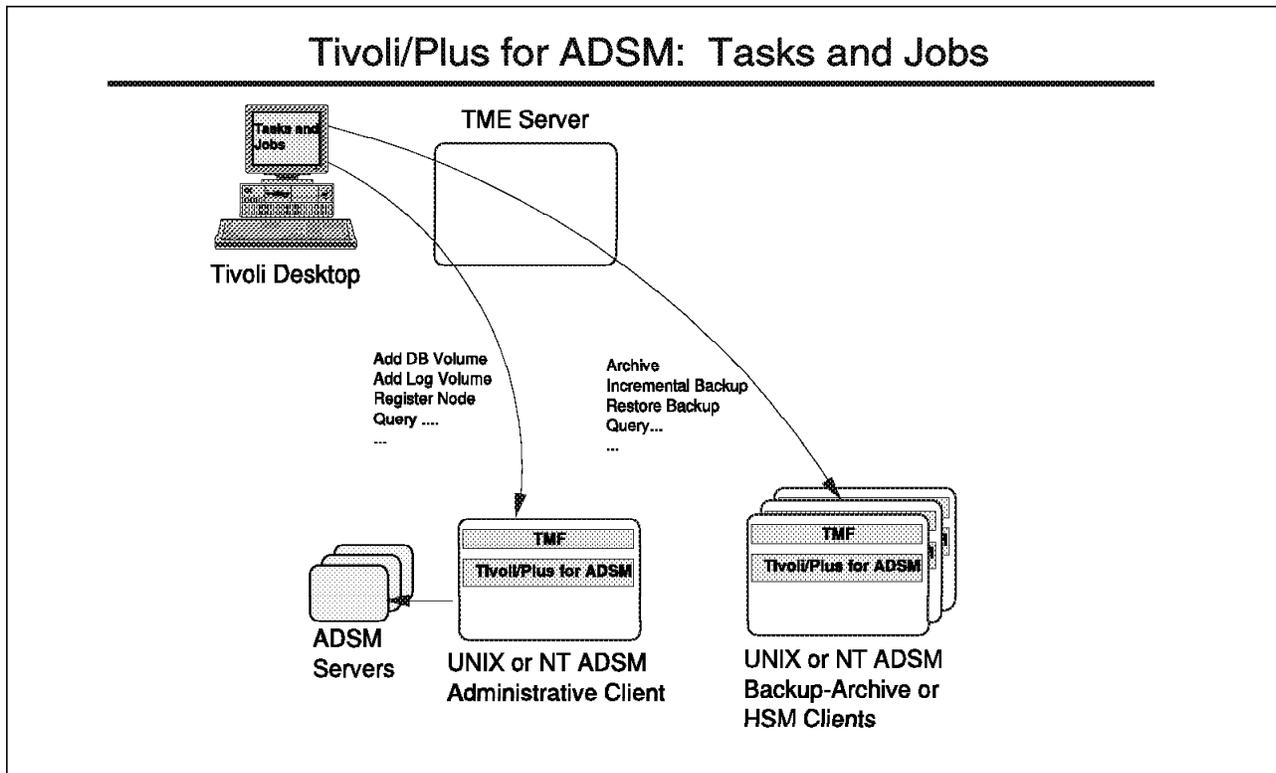
Tivoli/Plus for ADSM: Tasks and Jobs

When the Tivoli/Plus for ADSM module is installed, the Tivoli/Plus for ADSM collection is added to the Tivoli desktop. The Tivoli/Plus for ADSM collection contains icons that are called tasks and jobs. These icons enable you to perform functions on multiple ADSM servers and clients in the TME environment.

Tasks and Jobs

A task or job can consist of multiple ADSM commands. For example, the Add Log Volume job performs a DSMFMT, DEFINE LOGVOLUME, and then EXTEND LOG. The important point is that you can run a job on multiple servers or clients. The output from the job can be directed to your desktop or saved to a file. You can also specify whether a job runs serially on each machine, in parallel on all machines, or staged in groups of machines.

Some examples of ADSM server tasks and jobs include:

• Add DB Volume
• Add Log Volume
• Audit Licenses
• Define Association
• Register Node
• Query Events
• Query Volume

Some examples of ADSM client tasks and jobs include:

• Archive
• Incremental Backup

- Selective Backup
- Restore Backup
- Retrieve Archive
- Query Backup

Tivoli/Plus for ADSM: Benefits

The most fundamental service that the Tivoli/Plus for ADSM module and TME 10 delivers into an ADSM environment is improving the level of availability of the systems and network on which ADSM depends. Higher availability of these underlying resources through actions such as event monitoring and automated operations helps ADSM back up data quickly, limiting the chance that the only copy of the data will be damaged. Greater infrastructure availability means that data can be restored immediately, limiting application downtime and resultant revenue loss.

Automated Distribution of ADSM Client Code

Since ADSM on most platforms requires an interactive and therefore attended installation process, the software distribution function of the TMF framework can be further customized to support the unattended installation of ADSM. Courier can be used to distribute and then execute a script that mimics an attended installation process. Updating ADSM through PTFs can be accomplished in the same way. (The Tivoli/Plus for ADSM module does not automatically provide this function, but you can tailor courier to distribute PTFs.)

Centralized Launch of ADSM Servers

Without TME and the Tivoli/Plus for ADSM module, there is a one-to-one relationship between each ADSM server and its administrative console. With TME and the Tivoli/Plus for ADSM module, administrative control is based on a consolidated view of the resources managed, and multiple ADSM server management processes are brought together. This consolidation of events uses filtering and suppresses expected messages. Automated responses generated by TME are also automatically sent to the target ADSM server.

Centralized Launch of ADSM Administrative and Client Commands

ADSM administrators can issue commands to multiple ADSM servers or clients through TME. For example, you can change the backup schedule for all sales offices (each with its own ADSM server) from weekly to daily with only one command. This centralized launching greatly improves your productivity.

Centralized Monitoring of ADSM Status and Events

Without TME and the Tivoli/Plus for ADSM module, there is a one-to-one relationship between each ADSM server and its server console. With Distributed Monitoring, T/EC, and the Tivoli/Plus for ADSM module you have a central place that can receive ADSM server and client events and messaged. Distributed Monitoring also proactively checks for events that do not send a message.

Coordination of ADSM Operations with Other Applications

In addition to proactive monitoring and automation, you can use TME to coordinate the action of other management applications with ADSM to improve data reliability. For example, integrating the actions of the TME 10 central job scheduling facility with ADSM ensures that a database has been updated before it is backed up.

No Further Education Needed for the ADSM End User

TME enables you to execute all end-users tasks, such as backing up and restoring data, from a remote location. Eliminating ADSM end-user training can be a great advantage if end users do not have the appropriate skills or the value of their time is high.

Tivoli/Plus for ADSM: Requirements

Supported Platforms

- ADSM Admin, Backup-Archive, or HSM Clients
  - AIX 3.2.5 and 4.1
  - HP-UX 9.x and 10.0
  - NCR SVR4
  - Sun OS 4.1.3
  - Sun Solaris 2.3 or 2.4
  - Windows NT 3.5 or 3.5.1

- Software Distribution Only
  - DOS
  - NetWare 3.1, 3.12, or 4.0
  - OS/2 2.x
  - Windows 3.11

Software Requirements

- Required
  - ADSM Code
  - TME 10 Management Framework
  - TME 10 Enterprise Console
- Recommended
  - TME 10 Distributed Monitoring
  - TME 10 Software Distribution

Hardware Requirements

- 0.5 MB disk space (programs)
- 50 KB disk space (messages)

Tivoli/Plus for ADSM: Requirements

These are the requirements for the Tivoli/Plus for ADSM module:

Supported Platforms

The Tivoli/Plus for ADSM module is supported on the following platforms:

ADSM Administrative, Backup-Archive, or HSM Clients

- AIX 3.2.5 and 4.1

- HP-UX 9.x and 10.0

- NCR SVR4 MP-RAS 2.03

- SunOS 4.1.3

- Sun Solaris 2.3

- Windows NT 3.5 or 3.5.1

Using the Tivoli/Plus for ADSM module for clients is optional. You could just use Tivoli/Plus for ADSM to manage and monitor ADSM servers and still have a lot of function available to you.

Software Distribution Only

The Tivoli/Plus for ADSM module can be used to distribute software to the following additional clients:

- DOS

- NetWare 3.11, 3.12, or 4.0

- OS/2 2.x

- Windows 3.11

Software Requirements

The Tivoli/Plus for ADSM module has features that use TME 10 Software Distribution and TME 10 Distributed Monitoring. These applications must be installed, otherwise some specific features of the Tivoli/Plus for ADSM module may not work.

When installing Tivoli/Plus for ADSM module, ensure that the following products are planned for installation on the TME 10 managed node:

- ADSM administrative client, backup-archive, or HSM client code
- TME 10 TMF Management Framework (required)
- TME 10 Tivoli/Enterprise Console (recommended on the TME Server)
- TME 10 Distributed Monitoring client (recommended)
- TME 10 Software Distribution (recommended)

To use distributed monitoring, message consolidation, and event filtering for ADSM clients, the ADSM client must also be a TME managed node. If the ADSM client is not a TME managed node, you can still perform software distribution and you can monitor and manage your ADSM servers.

Hardware Requirements

The programs and the messages on each machine where the Tivoli/Plus for ADSM module is installed require. This, respectively, 0.5 MB and 50KB. This space is in addition to the space requirements for the management platform, the Tivoli application, and the ADSM storage software.

More Information

For more information about TME 10 and Tivoli/Plus for ADSM see:

- *TME 10 Tivoli/Plus for ADSM User's Guide* (GC31-8405)
- Tivoli announcement letter ZP96-0411 (EMEA) 296-216 (US)
- *TME 10 Cookbook for AIX: Systems Management and Networking Applications* (SG24-4867)
- Tivoli TME 10: A New Era in Network Computing Environment (TME10OVS TERS3820 on MKTTOOLS)

Chapter 4. Andrew File System and Distributed File System Support

AFS and DFS: Description

- Distributed File Systems
- AFS
- AFS Support
- DFS
- DFS Support
- DFS and AFS
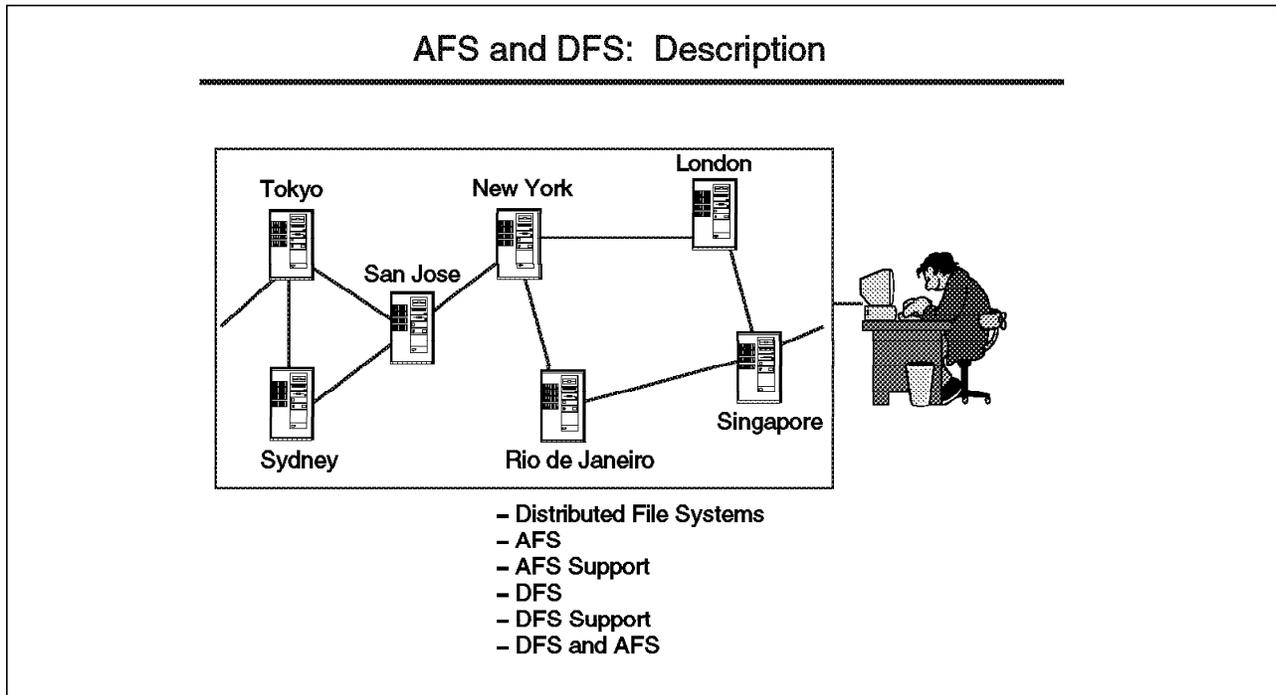
AFS and DFS: Description

In an ideal distributed environment, different computers, operating systems, and networks should interact as if they were all part of a single system. The Distributed Computing Environment (DCE) from the Open Software Foundation (OSF) is a set of services that can help achieve this goal: to deliver information from wherever it is stored to wherever it is needed, without exposing the network complexity to the users.

Distributed File Systems

A distributed file system provides the ability to see a single image of a file system even though parts of it (that is, individual files or directories) may be scattered across many machines. The file system has a single set of file names. As the user, you do not have to know where a file is located; you only have to know the path to it in the directory structure. The way you retrieve the information (or file) should be the same from any location, and ideally the locations can be anywhere around the world. You want to have an infinite workspace made up of all information in the global file system, and without restrictions defined by specific organizations, locations, or machines. You also want to address the files by standard office applications on your desktop system.

Although information must be easily shared within and between organizations, there must be ways of controlling access to the information—not only who can get into the network but also which files each individual can access.

Andrew File System

The AFS is a good example of a globally distributed file system. Developed by and available through Transarc Corp., AFS is based on a distributed file system originally developed at the Information Technology Center at Carnegie-Mellon University. However, AFS is not recognized as a UNIX standard.

AFS establishes a single shared file space for all authorized user from all possible locations. It is made up of cells, where each cell is a collection of servers grouped together administratively to present a single cohesive file system. Typically an AFS cell is a set of hosts using the same Internet domain name. An AFS cell can be from one to thousands of units.

AFS makes use of the Kerberos authentication system to verify the users and access control list (ACLs) to control directory access. Passwords are never sent across the network in plain text, and visible encrypted password files are not necessary. There is also mutual authentication: both the service provider and the service requester must prove their identities.

In AFS, the cache manager keeps a local cache of all file accesses to remote systems. Subsequent accesses then finds the data locally, providing high response and keeping the network traffic down. Workstations without disks have the data cached in memory.

The AFS server keeps the file mapping between the directory name and location, making the file space location independent. Location independence means that a user only has to know the path name of a file, not where it is located. With location independence comes scalability. New clients and servers are easily added as there is no need to make any changes in the existing systems.

Files in AFS are stored in structures called *volumes*, which reside on the server. Volumes containing frequently accessed data can be read-only replicated on several servers, thereby providing better response and high availability. If one server goes down, the data is automatically accessed from another replica.

The communication protocol is based on TCP/IP and is optimized for wide area networks (WANs) and for as little retransmission as possible. There is also interoperability with network file system (NFS), which is especially useful in a transition period from NFS to AFS networking. This function is provided by the NFS/AFS translator. The translator also provides an AFS solution for systems that are not directly supported with a full AFS system.

Andrew File System Support

AFS supports a number of UNIX workstations, such as:

- Digital DECstation and Alpha
- HP 9000 Series 700 and 800
- IBM RS/6000
- SGI R3000 and R4000
- Sun Sun4, SPARCstation, SPARC10, 600MP, and SPARCcenter

  and versions exist for:
- CONVEX C-Series
- MVS/ESA
- Control Data CD4000

Other "as is" ports also exist.

AFS and DFS are similar in architecture and offer many of the same benefits to users. DCE-oriented sites can migrate from AFS to DCE/DFS with the AFS/DFS Migration Toolkit.

Distributed File System

As part of the DCE environment, the OSF has defined a standard for distributed file systems. Transarc Corp. has implemented the standard in its DSF product, which is the key information-sharing component of DCE. Thus DFS is designed to be POSIX compliant.

DFS joins the file systems of individual workstations into a single virtual file system with a single name space, and it provides a consistent interface for users anywhere in the environment. DFS is built on top of DCE and uses the services provided in DCE. For example, for security, DCE Security Services provide mechanisms for authentication, privilege service, and registry service. This information is used to determine a user's access rights to the DFS files and directories.

DFS provides enterprisewide file access and management for LANs and WANs, enabling many computers from many locations to operate as a single unit.

DFS is available from Transarc Corp., and from other leading client/server vendors on all major platforms, from desktop PCs to UNIX systems and to MVS on the mainframe.

The DFS client includes a cache manager. When a file is read from the file server, it is cached locally for subsequent references. The server keeps track of the client cache, guaranteeing cache consistency. Files in DFS are grouped in file sets, which are hierarchical groups of files managed as a single unit. The concept of file sets is similar to the concept of volumes in AFS.

The files most often used are replicated onto multiple file servers. Replication spreads the load across several server machines, avoiding bottlenecks and overloaded systems. Through replication, high availability of data is achieved. If a client loses access to the current server, it automatically switches to an available replica.

A token manager on the server cooperates with the token management layer in the client and synchronizes access to the files by many clients simultaneously.

DFS maintains strict levels of security through encrypted login mechanisms using Kerberos and highly tunable ACLs for the files and directories.

For further information, refer to the redbook entitled *The DFS for AIX/6000* (GG24-4255).

Distributed File System Support

DFS is available for UNIX and other workstations that support DCE, and DCE is a prerequisite on the workstation. DFS is also available in a scaled down version for Windows 95, Windows NT, and Windows 3.11 that does not require DCE on the desktop.

Distributed File System and Andrew File System

DFS was developed on the basis of experience with AFS, and in many ways it can be considered an upgrade of AFS, but with some significant differences:

- In DFS, files on machines with different operating systems are available in the single distributed file system. For example, AIX and OS/2 users can share files with DFS without having to copy them back and forth between the systems.

- With DFS, ACLs are set for individual files. In AFS the ACLs are set for the directory level. In DFS users are allowed to set default "inheritance ACLs" for subdirectories or for files created from a directory owned by the user.

- DFS is fully integrated with the other DCE services, namely the DCE Security Service, Directory Service, and Distributed Time Service.

- DFS supports the POSIX single-site semantics. If a process running on one client writes to a file, all other DFS clients see that change on the next access to the file.

- DFS is integrated with AIX permission bits. There is also support for exporting local AIX JFSs and maintaining access to them through the local mount path.

AFS and DFS: Previous Backup Support

- ADSM File-Level Backup
  - ★ No Volume Dump
  - ★ Limited ACL Support
- butc
  - ★ No File-Level Backup
  - ★ Preserves Attributes
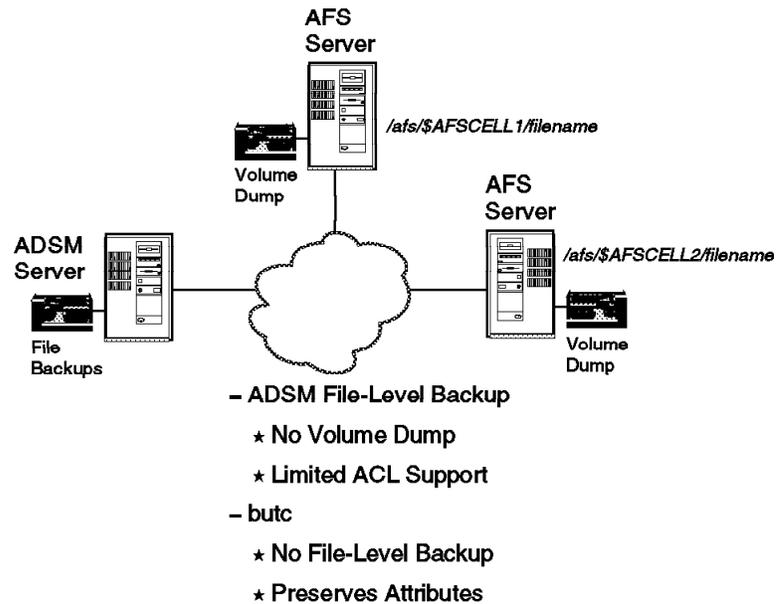
AFS and DFS Previous Backup Support

Backup requirements vary, including single files, full disks, and AFS and DFS volumes. With a global, single-view file system, only the actual local data can be backed up. The support for AFS or DFS backup in ADSM is limited.

ADSM File-Level Backup

The ADSM clients allow AFS and DFS files to be backed up just like any other files in the hierarchical file structure. The AFS ACLs can be backed up only by the AIX client, and the DFS ACLs cannot be backed up from any platform.

There is no support for backing up full volumes or file sets. A volume can contain mount points of any other volume, and volumes are logically connected through mount points to form a tree of directories and files. The ADSM clients cannot back up volumes with their attributes or mount points. Furthermore, single files cannot be backed up at an acceptable speed in large cells.

butc

Transarc's AFS backup commands back up AFS volumes to tape devices, using a program called butc (BackUp Tape Coordinator). Each machine running butc is connected to a tape device, which can be a drive or a stacker. butc is the most widely used AFS and DFS backup program to date. It is a volume backup system used to dump volume images to tape devices. The dumps preserve not only files and directories but also ACLs, volume attributes, and mount points. However, the minimum unit of dump and recovery is a volume; that is, butc provides no support for file-level backup and recovery.

AFS and DFS: New Backup Support

- Volume Dump Support
  - ★ buta
- File Backup Support
  - ★ Mount Point Backup
  - ★ DFS Include/Exclude
- AIX Clients Only
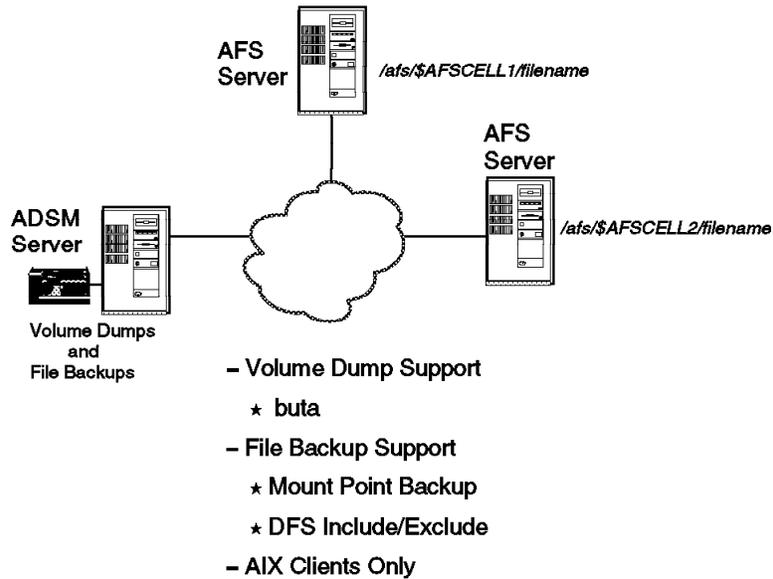
AFS and DFS File Systems: New Backup Support

New backup support has been added for the AIX client. This new function improves ADSM support of AFS and DFS. Specifically, it adds support for volume backup, which includes mount points and ACLs. It also enhances file backup support to include mount points and ACLs. It further adds some functions to address the management of distributed file system backup. These functions include options for excluding directories and file spaces from backup, and environment-variable support for user-specified virtual mount points. All of these new functions are available on the AIX platform only. For more details about option settings and how to use the AFS and DFS backup commands, see *ADSM AFS/DFS Backup Clients V2.1* (SH26-4048).

The new AFS and DFS backup functions consist of four separate programs, all available for AIX 4.1.4. These programs are called through the AFS and DFS backup commands issued by the backup administrator.

- AFS volume backup agent (buta - BackUp to ADSM)

  To improve and centralize the backup process and enable the use of the ADSM server instead of tape devices, buta functions as an ADSM interface program between the AFS backup commands and an ADSM server. The buta program operates with the AFS 3.4a buserver, backup database, and backup commands. Note that you can use the 3.4 version of the backup system in a 3.3x cell. However, all backup programs (buta, buserver, backup) should be on the same version.

- DFS file set backup agent

  The DFS file set backup agent program is the DFS counterpart of the AFS buta program. It works with AIX DFS 2.1.

- AFS file backup client

  The AFS file backup client program enhances the AFS backup support to back up both ACLs and mount points. The previous ADSM 2.1 client could only back up ACLs.

- DFS file backup client

  The DFS file backup client program enhances the DFS backup support to back up ACLs and mount points.

Volume Dump Support

The AFS and DFS buta command is similar to butc but uses ADSM client API calls. Through the API, each dumped volume is sent to the ADSM server as a file. The volume dumps can be incremental or full. The file name is the same as the name of the volume. The volume dump files associated with a dump are stored in a single file space on the ADSM server, and the name of the file space is the dump-id string. Several buta backup streams from different file servers can be active concurrently.

Dumps can be deleted from the ADSM server by deleting the corresponding file space. This can be done within an ADSM administrator session. Deleted dumps should also be removed from the AFS or DFS backup database through the AFS or DFS backup command. A utility program, delbuta, deletes the unwanted old dumps from both the AFS and DFS backup database and ADSM server at the same time.

There is no national language support for AFS and DFS backups. These programs are designed to be an integral part of the AFS and DFS volume backup system, which does not address national language support.

File Backup Support

The new file backup support with the ADSM AIX client enables you to back up or restore AFS and DFS files and directories. In addition you can back up ACLs for directories as well as mount points.

Several new options are introduced:

Mount Point Backup

Because users can mount the same volume at more than one point in the file tree, it is necessary to control under which path files are backed up.

With AFSBACKUPMNTPNT or DFSBACKUPMNTPNT set to yes, directories are checked to see whether they are mount points. If they are, they are backed up only as mount points (symbolic link) and not viewed as directories.

DFS Include/Exclude

In the ADSM UNIX environment, client files are included or excluded from backup by being listed in an inclexcl file. In this file the prefix /... (3 dots) is interpreted as the arbitrary directory prefix.

The DFS include/exclude file is special. Because the prefix /... (3 dots) indicates the global root in DFS, a different include/exclude option (DFSINCLEXCL) is used to specify a separate DFS inclexcl file. In this file, /... (3 dots) is interpreted as the global root of DFS, and a /.... (4 dots) prefix is interpreted as the arbitrary

prefix. For a DFS file backup, the contents of the DFS inclexcl file is used to match the files for the include/exclude management-class actions.

Exclude File Spaces

The EXCLUDE.DIR.BACKUP option is used to exclude an entire directory from ADSM backup services. When you exclude a directory, you can still back up specific files within that directory by using selective backup.

The EXCLUDE.FS option excludes an entire file space from incremental backup. For example, select this option to prevent users from starting an incremental backup from the root of AFS or DFS.

Virtual Mount Points

The VIRTUALMOUNTPOINT option defines a unique virtual mount point for a file system with files that you want ADSM to back up.

Domains

The DOMAIN option specifies the file systems you want to include for incremental backup in your client domain. You can include virtual mount points in a domain.

AIX Clients Only

The new AFS and DFS backup support is available for the AIX ADSM client only. You must either have the AIX client installed on your AFS server or on an AFS client somewhere in your network. Be aware that if your AIX AFS client is in a separate machine from your AFS server, the backups will increase your network traffic.

The AFS and DFS backup support comes with client level PTF 6 (2.1.0.6) and requires the ADSM chargeable feature, Open Systems Environment Support (previously called *OpenEdition Support*).

The software requirements are:

- AIX Version 4.1.4
- AFS Version 3.4.a
- DFS Version 2.1

Chapter 5. WebShell Client
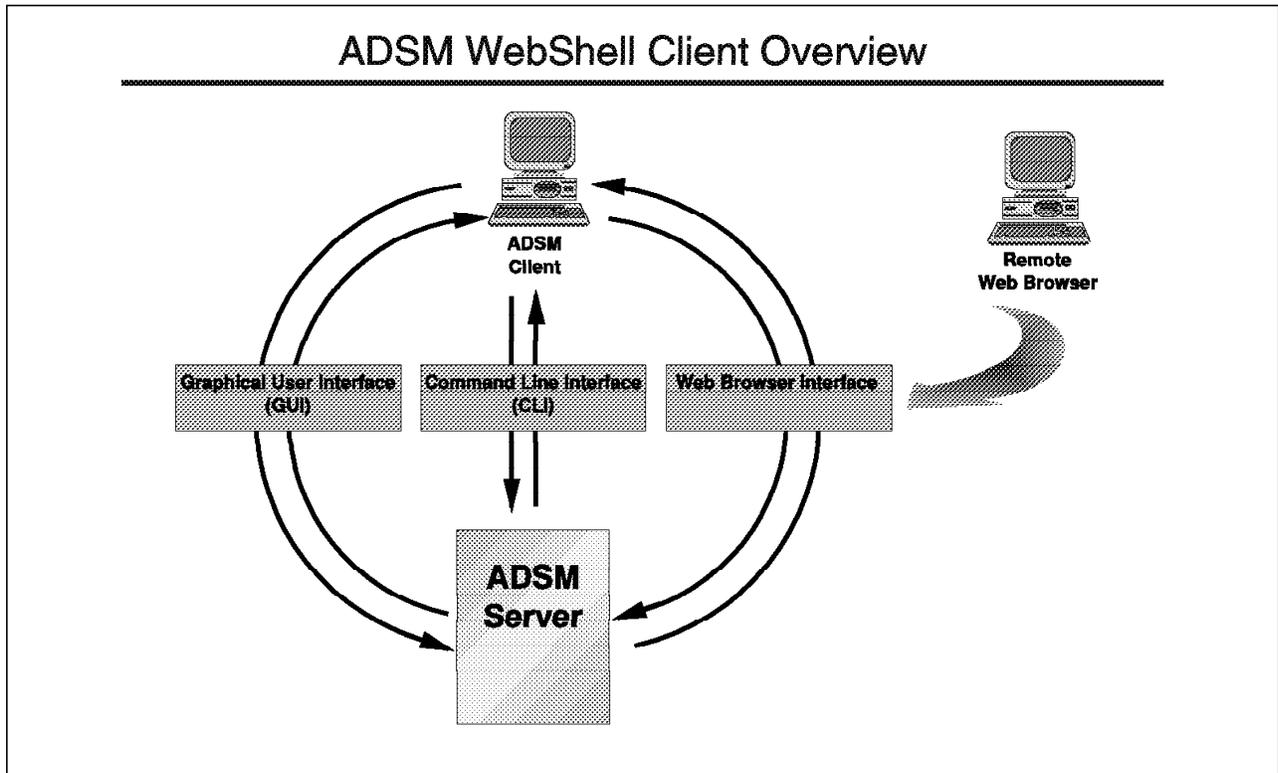
ADSM WebShell Client Overview

The ADSM WebShell client is a new interface to the ADSM backup-archive client. This Java-enabled interface is designed to provide end users with easy, mobile, backup and restore capabilities and help desk personnel with remote backup and restore capabilities.

The ADSM WebShell interface enables a web browser from any platform and located anywhere on the network to interface with ADSM. The web browser talks to the WebShell interface on an ADSM client, which then talks to the ADSM server. The web browser does not talk directly to the ADSM server.

The WebShell interface is available for OS/2 Warp, AIX, Windows 95, and Windows NT clients. The WebShell client code is free and can be downloaded from the ADSM home page, http://www.storage.ibm.com/storage/software/adsm, or from the ftp anonymous site, index.storsys.ibm.com under /adsm/fixes/v2r1/os2. A demo is also available from the ADSM home page.

Webshell Client Principles
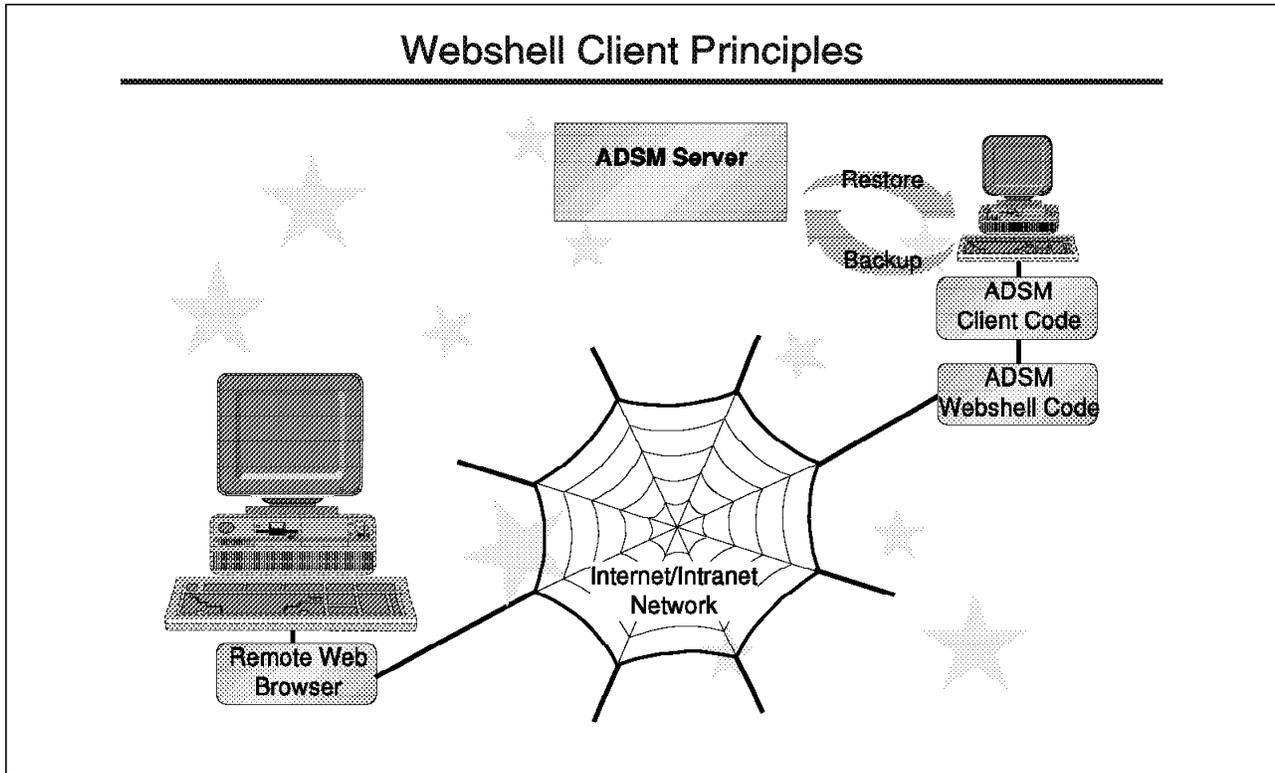
WebShell Client Principles

The ADSM WebShell client creates a mini-web server on the ADSM
backup/archive client and therefore makes it accessible through the Internet just
like any other uniform resource locator (URL). Once the WebShell code is
downloaded, it has to be started either in a window (OS/2, NT) or as a
background process (AIX). The Internet address is the TCP/IP address of the
ADSM client (cerium.almaden.ibm.com in this example). The default port is
2121, and the directory is webshell.

To connect to this new URL, point your web browser to the following address:

- **http://cerium.almaden.ibm.com:2121/webshell**

WebShell Client Security


```
┌──────────┐  WebShell   ┌──────────┐   ADSM    ┌──────────┐
│   Web    │ ◄────────►  │  ADSM    │ ◄────────►│  ADSM    │
│ Browser  │             │  Client  │           │  Server  │
│          │  Security   │          │  Security │          │
└──────────┘             └──────────┘           └──────────┘
```
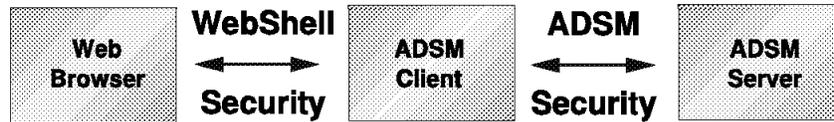
- Separate WebShell Userid and Password
  - Protects against Unauthorized Use
  - ADSM Client Password Not Transmitted between Browser
    and ADSM Client

WebShell Client Security

The WebShell client allows you to operate in an Internet or Intranet environment. It is important to the information stored in the ADSM Server and the integrity of the client workstation that the WebShell client operates in a secure environment.

Separate WebShell Userid and Password

A user accessing the WebShell client from a browser has his or her own userid and password. Remember that multiple users can access the same ADSM client, so it is possible to have more than one WebShell userid.

The WebShell code on the ADSM client workstation uses the ADSM node name and password to communicate with the server. The ADSM node name and password are defined to the WebShell code during installation. For more information about ADSM authentication, refer to "Authentication" on page 30.

WebShell Client Interface
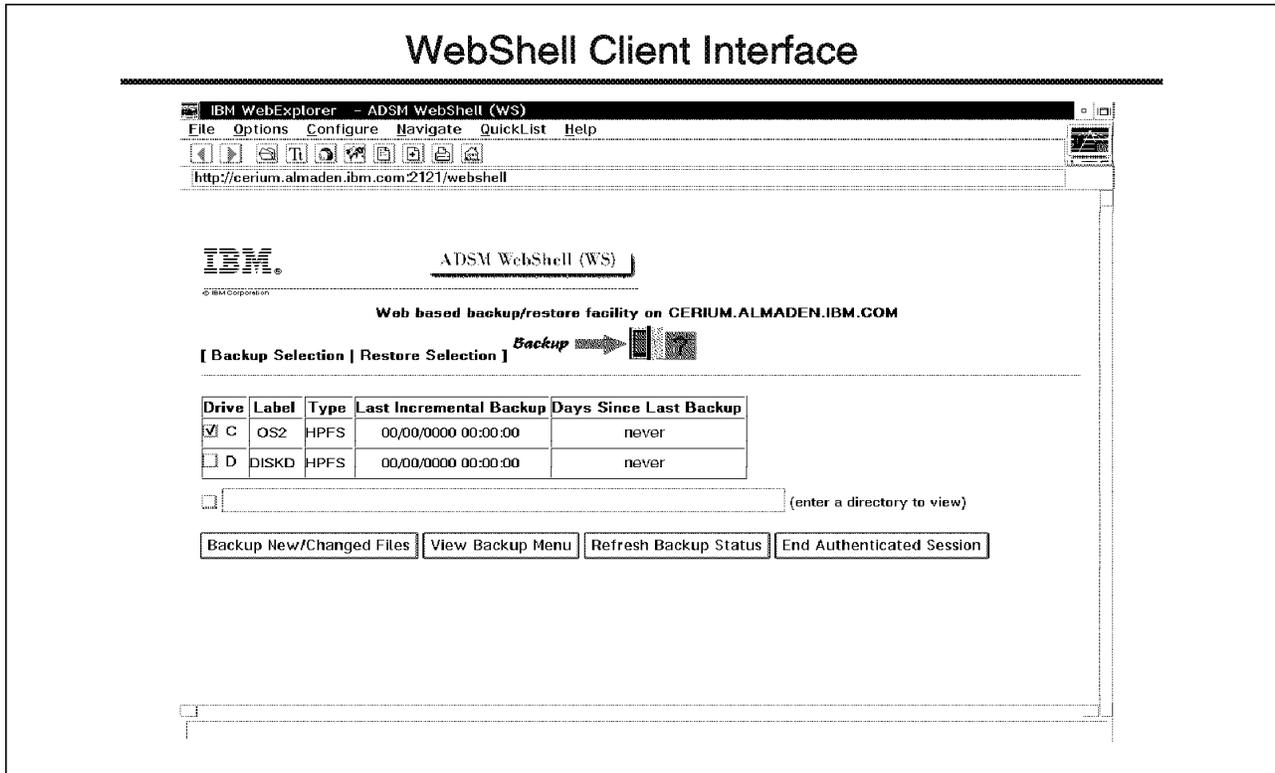
WebShell Client Interface

The ADSM WebShell client offers a typical web interface with HTML tags and push buttons, thereby facilitating the use of ADSM for those who are not familiar with it.

Through the web browser, you can select the following client operations:

- Display status of backup operations

- Incremental and selective backup

- Restore of the last saved version of a file (=active version)

The selection is converted to an ADSM backup client command passed to the ADSM server. The results of the command are then converted to HTML tags and displayed through the web browser.

WebShell Client Benefits

- Remote and Central Control of LAN and File Servers
- Administrator/Help Desk Productivity
  - Consistent Interface across Multiple Platforms
- End-User Productivity
  - Familiar, Consistent Interface
  - Task-Switching Effort Reduced
- Lower Costs
  - Reduced Training Requirements

WebShell Client Benefits

The ADSM WebShell client provides these benefits:

Remote and Central Control of LAN File Servers

You no longer have to be physically on the same machine as the ADSM client in order to use it. Remote operations can be done from a central location.

Administrator/Help Desk Productivity

Help desk personnel and administrators can use a single web browser interface with which they are already familiar to control multiple ADSM clients on various platforms.

End-User Productivity

End users can use their familiar web browser to restore and manage files.

Lower Costs

Using a familiar web browser to access ADSM clients reduces the training that would otherwise be necessary.

Chapter 6. NetTAPE

NetTAPE Overview

Tape Management ?

Drive Management ?

Library Management ?

NetTAPE ?

ADSM ?

Drive sharing ?
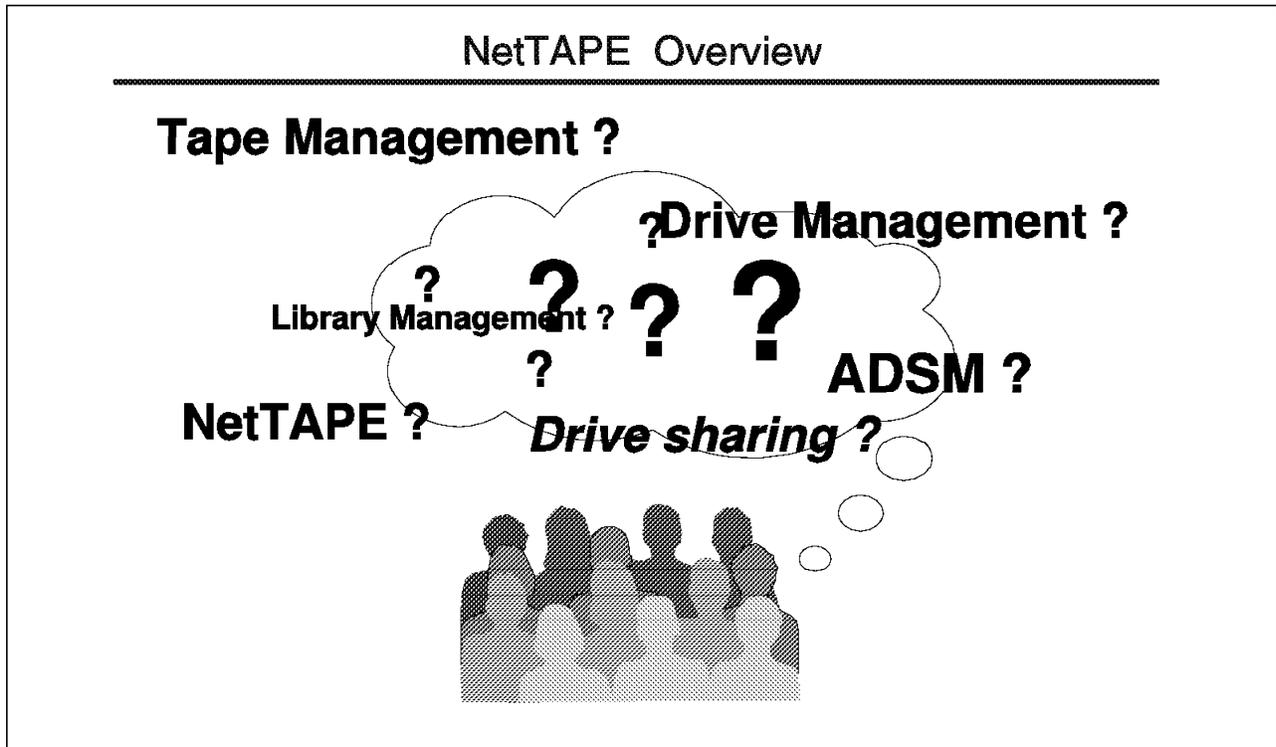
NetTAPE Overview

Tape management enables you to locate tape volumes, label them, register expiration dates, and issue some recycling commands when necessary. A tape management product may also include the ability to segregate tapes according to the application or software that allocated them or the machine to which they belong. Tape management can also include drive management, by sharing drives and tape libraries among several hosts.

AIX by itself does not provide any of these tape management functions. ADSM provides some, and NetTAPE provides most of these functions. Therefore, the use of NetTAPE in association with ADSM can be of great help, especially if the drives and tape volumes have to be segregated according to the needs of the applications.

NetTAPE is a collective name for two program products, IBM Network Tape Access and Control System for AIX and IBM NetTAPE Tape Library Connection, which contains the automated library support. NetTAPE enables access to tape devices and consolidated tape operations on a network of AIX workstations, providing a single-system image to users. NetTAPE supports both RS/6000 workstations and RS/6000 Scalable POWERparallel Systems (SPs).

Tape Management with AIX

AIX by itself provides no tape management functions. There is no catalog for locating data, no way of sharing drives, and no way of protecting a tape from being overwritten by any user who issues a write-on-tape command (if you forget to remove your tape from the drive he or she is accessing).

When a tape drive is available, someone can walk up to the drive, insert a cartridge, and issue some AIX commands that will read or write data on the cartridge. This process is fine if you have a tape drive for your own private use, but it is far from sufficient if several tape drives are used by different users and applications. If a drive is not available, and you try to use it, you get an I/O error.

You have to install additional products to use tape drives and libraries in a more automated production environment. ADSM and NetTAPE are two such products.
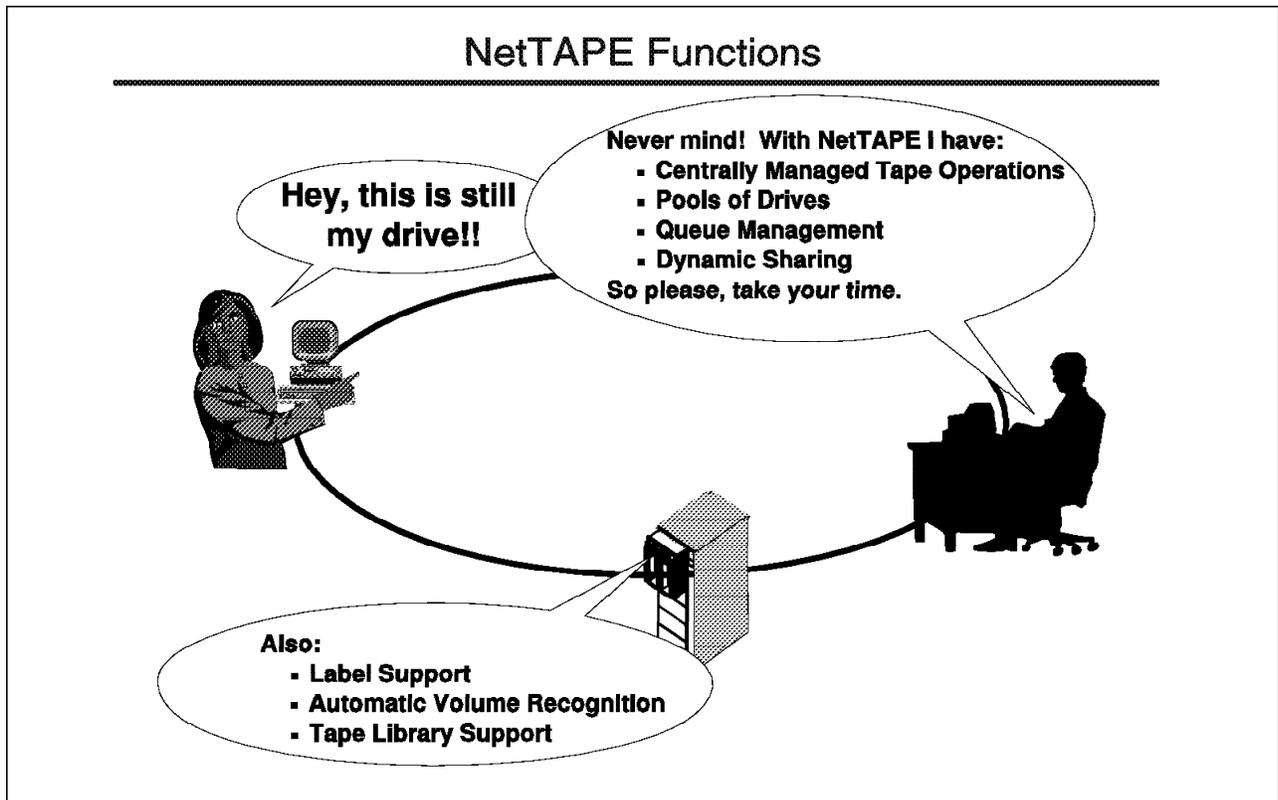
NetTAPE Functions

The main objective of NetTAPE is to provide a single-system view of tape operations for the user or application.

Centrally Managed Tape Operations

NetTAPE provides facilities for centrally managing tape operations. Tape operators, from one or several operations stations, get mount and unmount messages for a set of tape devices attached to multiple systems. Tape operations are typically done under the NetTAPE XWindows, Motif-based GUI. A CLI is also available but does not provide all the functions the GUI offers. Users can ask for a drive by using a generic name, which is a device pool containing several drives.

Pools of Drives

NetTAPE treats all of the drives and libraries as global resource pools, thus allowing better allocation strategies than if device selection were left to the user. With NetTAPE users can access any tape drive in the NetTAPE network.

Queue Management

NetTAPE processes all mount requests issued by the NetTAPE clients, manages the queues for devices, but also enables users to change the order of a mount request in the queue for a certain device or delete a queued mount request.

Dynamic Sharing

NetTAPE can be used to share tape devices dynamically among multiple AIX applications, including ADSM for AIX.

NetTAPE supports the tape label formats usually used on mainframe systems such as MVS, thus enabling the interchange of tape cartridges between MVS and AIX applications.

Label Support

NetTAPE supports tape files with record formats and tape labels typically used on mainframes but not normally supported on UNIX. Thus users can process tape volumes created by mainframe applications in an AIX environment.

Automatic Volume Recognition

NetTAPE provides automatic volume recognition, which allows you to premount volumes that you know an application will use later.

Tape Library Support

NetTAPE provides support for automated tape libraries, thus integrating tape library operations with manual drives.

NetTAPE does not contain a tape catalog function or any link to an external tape management system catalog, but it provides a facility for recording basic volume information. The information is stored in volume files, with one volume file per tape file or data set.
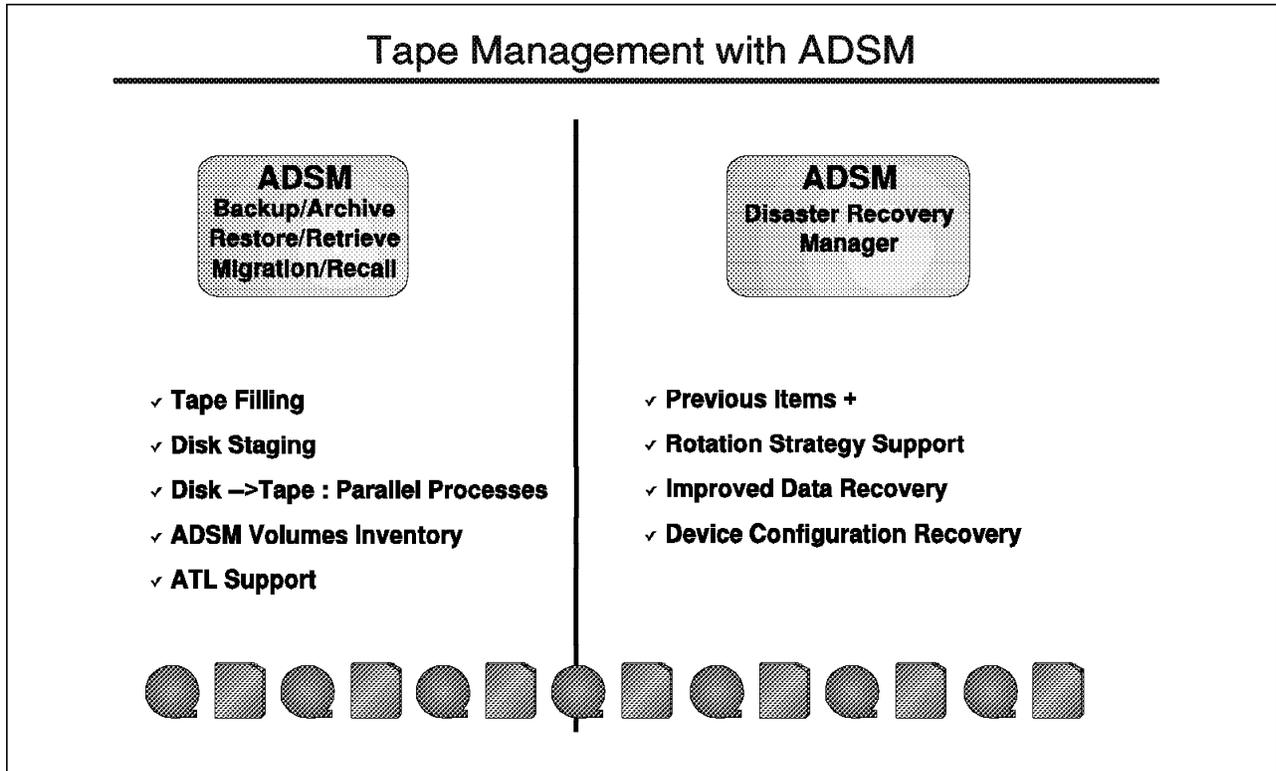
Tape Management with ADSM

ADSM provides tape management functions that manage ADSM tape volumes, but it cannot manage tape volumes for other applications. The ADSM server is responsible for managing the ADSM storage volumes, and it uses the database to keep track of the location of every piece of data.

DRM brings some specific enhancements to tape management support, as detailed here.

ADSM Basic Functions

Using ADSM to support tapes enables you to get away from a model where a user owns a tape volume. When you use ADSM, user data is managed in a way that makes tape volume transparent to the end user.

Some of the benefits are:

- Tape filling: improved tape volume utilization

  ADSM tape volumes are filled to full capacity. When the data is deleted from these volumes, the space is reclaimed and the tape can be reused.

- Disk staging: improved tape drive utilization with disk storage pools

  The ability to have a disk storage pool in front of the tape storage pool avoids the problem of ADSM clients holding on to drives for many hours.

- Disk --> tape: parallel processes. During the process of migration from disk to tape or during the storage pool backup process, ADSM can launch several parallel processes, thus providing an increase in data transfer performance.

- ADSM volume inventory

  You do not have to manage the ADSM volume inventory by yourself. ADSM does it for you.

- ATL support: ADSM can manage the robotic moves inside an automated tape library, without the need of another support program. Thus it is easy to implement tape automation.

ADSM Disaster Recovery Manager

DRM enables you to reconstruct your ADSM server and clients at a recovery site. It provides such tape management functions as:

- Rotation strategy support

  DRM can manage an offsite rotation strategy for ADSM tapes.

- Device configuration recovery

  ADSM supports many different devices, which sometimes use apparently the same physical support but do not write the data with the same format (for example, 8 mm tapes can be 2.3 GB, 5 GB or 7 GB). By allowing you to back up your device configuration, ADSM helps you rebuild your environment at the offsite location by specifying which device, format, and volumes are required.

  Please refer to "Disaster Recovery Manager Overview" on page 73 for more information.

Remember that ADSM only manages its own storage pool volumes; it cannot manage volumes of other applications.

Tape Management with ADSM and NetTAPE
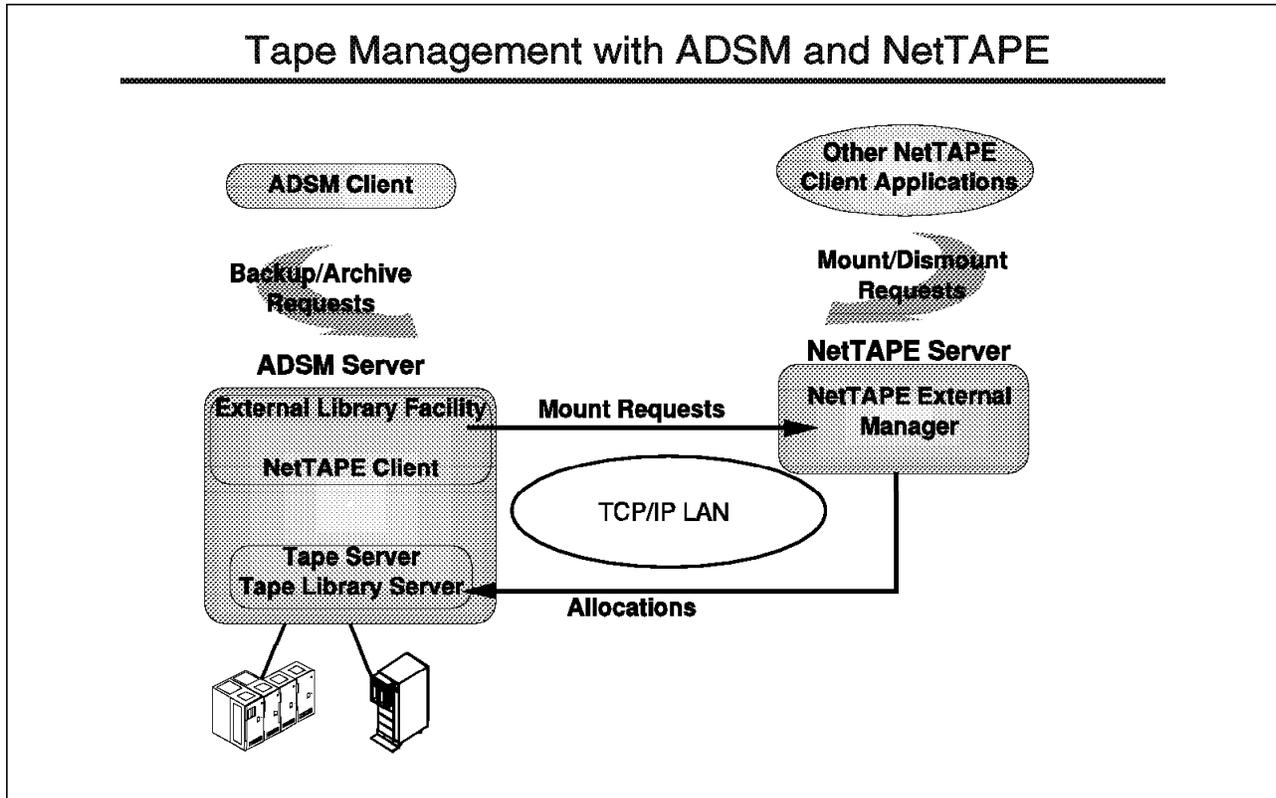
Tape Management with ADSM and NetTAPE

NetTAPE provides an interface to ADSM so that tape and library devices can be managed by NetTAPE instead of being dedicated to ADSM.

External Library Facility

The ADSM External Library facility is available for ADSM Version 2 for AIX. The External Library enables ADSM to decouple ADSM operations from the details of the tape drive and tape library configuration. When a library is defined as external, ADSM sends volume mounts, volume dismounts, and volume releases (returns to scratch) to an external process.

NetTAPE External Manager

The NetTAPE External Manager handles tape and library operations for ADSM external libraries. The libraries have to be defined to ADSM with the EXTERNAL parameter:

DEFINE LIBRARY library_name LIBTYPE=EXTERNAL
EXTERNALMANAGER=/usr/lpp/nettape/bin/cladsmd

ADSM sends a mount request to the NetTAPE External Manager, which is responsible for:

- Verifying that the request is coming from the ADSM server
- Sending a mount request to either a human operator or an automated library
- Handling eventual mount problems and, if possible, recovering from them

- Verifying that the correct volume is mounted

- Passing back to ADSM the device address where the tape volume has been mounted

ADSM opens the device, checks the tape label, and proceeds with writing or reading the tape volume.

The only things ADSM has to know about the tape device are its type and how many drives it can use at a time, as defined by the mountlimit parameter in the device class definition.

Facilitates Device Sharing

The NetTAPE External Manager enables you to share drives between ADSM and other applications.

By not dedicating drives to an ADSM server, the NetTAPE Library Manager selects any drive that happens to be free at the time ADSM requests a volume. The tape library and the tape drive that ADSM will use must be local devices. ADSM does not support remote access, even when going through NetTAPE.

Supports Tape Drives and Libraries

External libraries can be used with either automated or manual libraries. You can have several external libraries for different types of tape devices defined to an ADSM server. You can have some libraries defined as external libraries and others managed by ADSM.

NetTAPE provides full support for IBM 3420, 3480, 3490, and 3590 tape drives, plus the Ampex DST 310 tape drive. Support for 4 mm and 8 mm tape drives is limited to access only (no label, no volume recognition). The support of tape libraries is limited to IBM 3494 and 3495 tape libraries and STK 4400, 9310, 9311, and 9360 tape libraries.

More Information

For more information about tape management with ADSM and NetTAPE, see *AIX Tape Management*, SG24-4705.

Chapter 7. Network Storage Manager
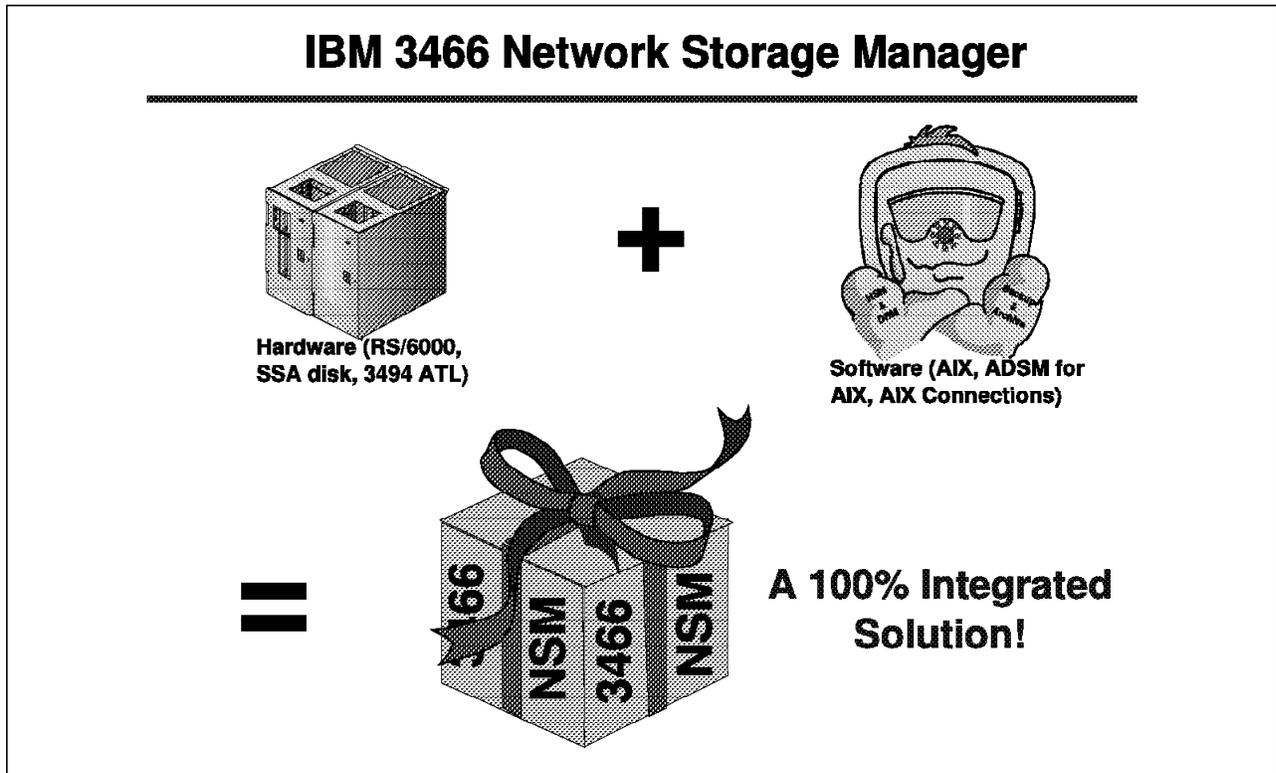
IBM 3466 Network Storage Manager

The IBM 3466 Network Storage Manager is a high-performance, open, storage management solution including hardware and software with ADSM functions.

The Network Storage Manager integrates proven software and hardware storage technologies into a single powerful solution for automated network backup. Both hardware and software have been integrated and tested for you.

Hardware and Software

The Network Storage Manager consists of:

- RS/6000 processor with a Serial Storage Architecture (SSA) disk

    The SSA disk can store from 72 to 144 GB of data.

- Network Connections

    Network Storage Manager can connect to FDDI, Token-Ring, or Ethernet networks. Depending on the configuration, up to four adapters of different protocols can be ordered.

- 3494 tape library with 3590 tape drives

    Depending on the model, you can have from two to four Magstar tape drives in the 3494 tape library. With a compaction factor of 3:1, the different models give you the following data capacity:

- Model B01 - two tape drives - 18.0 TB

  - Model B02 - three tape drives - 28.2 TB

  - Model B03 - four tape drives - 40.2 TB

  All tape cartridges are initialized and labeled when you receive them.

• Network Storage Manager Software Package

Network Storage Manager Software Package Version 1 Release 1 consists of:

• AIX Version 4 Release 1.4 with AIX Connections, including TPC/IP, APPC, NetBIOS, and IPX/SPX

• ADSM Version 2.1, including the WebShell client

• Communications Server for AIX Version 4

• AIX Internet Connection Secure Server Version 4 Release 1

• AIX REXX/6000

The software package is preloaded and configured. ADSM includes backup and archive functions, the DRM and HSM features, all current device modules, and any type of clients, except for the MVS OpenEdition client, which can be added separately at an additional charge.

```
┌─────────────────────────────────────────────────────────────────────┐
│                 IBM 3466 Network Storage Manager:                     │
│                            Benefits                                   │
│  ▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬  │
│                                                                       │
│                                                                       │
│     ➤ Complete, Integrated Solution                                   │
│     ➤ Powerful Storage Technology                                     │
│     ➤ Simple, Smart, and Automated                                    │
│     ➤ Cost-Effective                                                  │
│                                                                       │
└─────────────────────────────────────────────────────────────────────┘
```
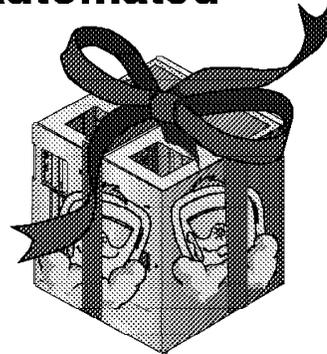
IBM 3466 Network Storage Manager: Benefits

Complete, Integrated Solution

The Network Storage Manager integrates leading IBM software and hardware technology into a single powerful solution. It is designed to help you protect your most valuable strategic assets—your business data.

Powerful Storage Technology

The Network Storage Manager includes some of the fastest and most reliable hardware technologies available today: SSA disks, 3590 drives, a 3494 tape library, and a RISC processor for power. It also comes with built-in powerful backup, archive, restore, HSM, and disaster recovery capabilities through the ADSM software.

Simple, Smart, and Automated

The Network Storage Manager is highly integrated and powerful, but it is also simple (preloaded, preconfigured GUIs), smart (web access), and automated (for example, scheduled operations and automated tape mounts).

Cost-Effective

From a system administrator point of view, less time has to be spent on installation and ongoing maintenance. Because all of the components have been preconfigured and pretested, there is no need to evaluate and test each piece of the solution. Therefore, you can better focus on other aspects of your network storage management job.

A StorageSmart Plug ′n Go Solution

When the Network Storage Manager is powered up, the server starts automatically with ADSM already configured.

Initial Configuration

The initial ADSM configuration includes:

- An administrative client called **admin**
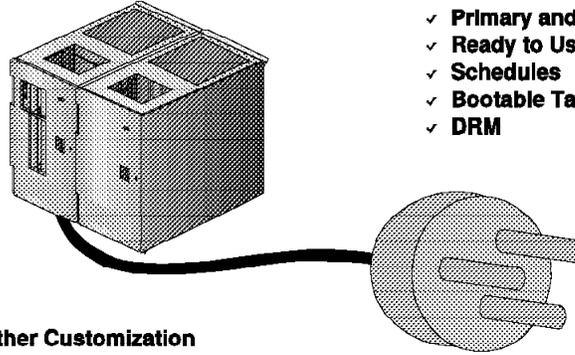
- Mirrored database and recovery log

- Primary storage pools for disk and tape and a copy storage pool for primary storage pools

- Tapes that are initialized, labeled, and checked in

- Schedules for backup of clients, inventory expiration, migration, backup of database and storage pools, and deletion of volume history entries

- A bootable tape containing a copy of the operating system, ADSM, and other associated software. In the event of a disaster or a disk crash, you can restore the system to its initial state from this tape. You then have to use DRM and ADSM to restore all other data to its current state.

- DRM has been configured, but there are some steps that you must take to prepare a disaster recovery plan file and track tapes that have been moved.

Further Customization

You can use the standard administrative client with the GUI or CLI to manage your ADSM server, or you can use the Web administrative interface provided with the Network Storage Manager. The Web administrative interface can access ADSM server functions from any workstation with a web browser that

supports HTML 2.0 or higher. The workstation does not have to have a standard administrative client installed.

Some of the functions you may want to customize are associating nodes with schedules and changing the password for **admin**.

Installing and Setting Up the Clients

- **Installing the Clients**
  - **– Download the ADSM Client Code from the**
  - **– Build the Options File**
  - **– Register the Node**
- **Using the WebShell Client**

Download ADSM Client Code

Restore / Retrieve

Backup / Archive

ADSM Client
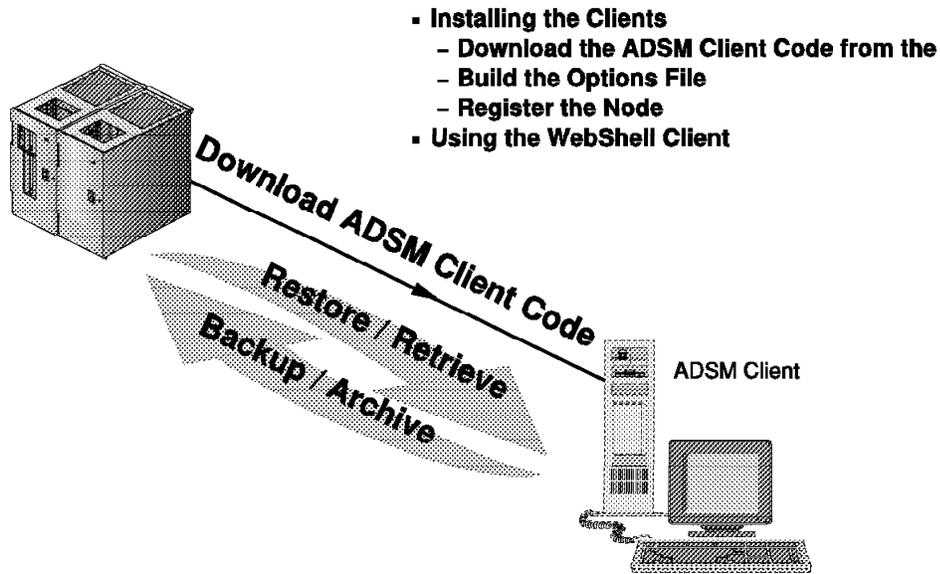
Installing and Setting Up the Clients

Before you can start using the Network Storage Manager, you have to set up at least one backup-archive client.

Installing the Clients

Once the Network Storage Manager has been installed, you download, install, and customize the ADSM clients by accessing the StorageSmart home page. The StorageSmart home page runs on the Network Storage Manager and allows you to:

- Download a copy of the ADSM client code

- Build a customized client options file

  The options file includes TCP/IP as the communication protocol and a default include/exclude list.

- Register your client with the server

Using the WebShell Client

The WebShell client is described in Chapter 5, "WebShell Client" on page 111. There are a few differences between the Web interfaces available with the Network Storage Manager and those available to other ADSM servers.

The Network Storage Manager StorageSmart home page has a front end to the WebShell client. The StorageSmart home page provides a list of ADSM clients on which the WebShell client is installed. You select a client from the list to access it.

Without the StorageSmart home page, you would have to know the ADSM client's TCP/IP address and type it in on your web browser as a URL.

The StorageSmart home page also provides a specific Web administrative interface, which is not available anywhere else. The Web administrative interface provides icons to easily query different resources on the server and allows you to enter ADSM administrative client commands.

More Information

For more information about the Network Storage Manager, see *IBM 3466 Network Storage Manager Introduction and Planning Guide* (GA32-0374) and *IBM 3466 Network Storage Manager Start Up Guide* (GA32-0373).

Chapter 8. Bare Metal Restore
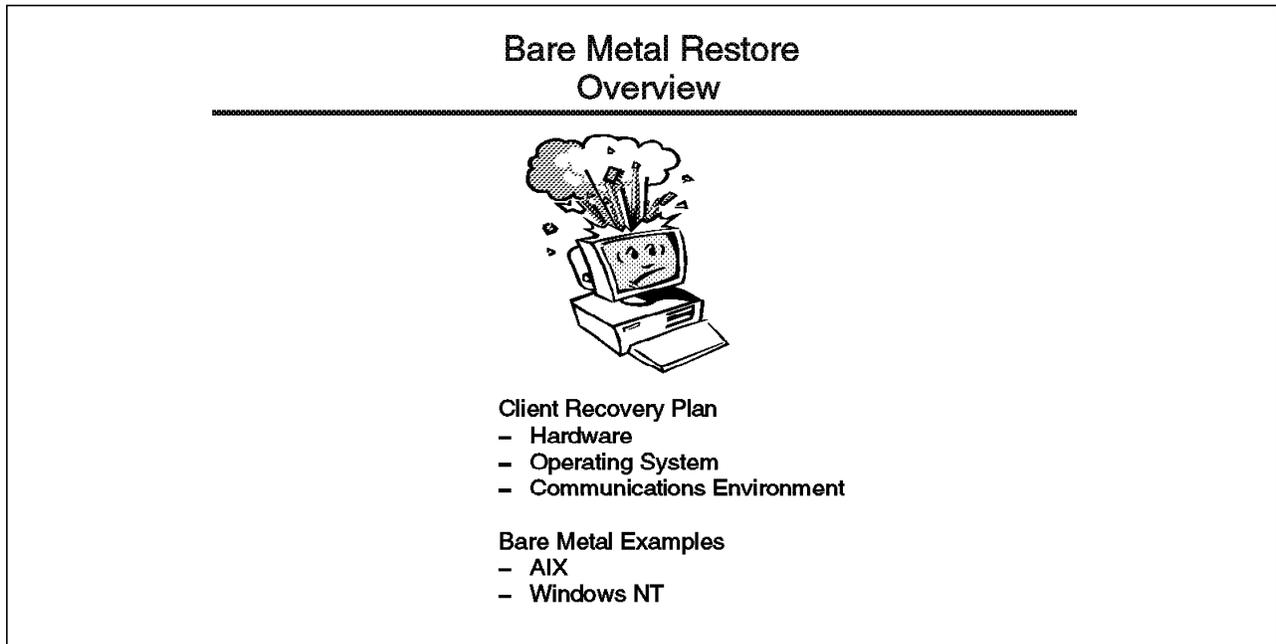
Bare Metal Restore: Overview

Bare metal restore is the recovery of an ADSM client from the "bare metal" up, including the operating system and communications software that has to be in place for the ADSM client to begin its restoration work. Bare metal restore includes using techniques for bootable recovery, recovery using remote distribution tools, and recovery using peer services to recover ADSM clients on the OS/2, Windows, Novell NetWare, and AIX platforms.

For further details about bare metal restore pf an ADSM client, refer to the redbook, *ADSM Client Disaster Recovery: Bare Metal Restore* (SG24-4880). Server disaster recovery for AIX is described in the redbook, *ADSM for AIX: Advanced Topics* (SG24-4601), and for MVS in *ADSM for MVS: Recovery and Disaster Recovery* (SG24-4537).

Client Recovery Plan

The recovery of ADSM clients depends on the existence of a re-covery plan to re-create the hardware, operating system, and communications environment necessary to run the ADSM client code. Once the basic operating functions have been restored, the client's file systems can be restored from the ADSM server storage pool.

The following items must be covered by such a recovery plan:

- Hardware, to re-create damaged units

- Operating system, to have a bootable backup

- Communications environment, to be able to communicate with the ADSM server or a system reload server

In addition it is necessary to know where to find the original installation software, and possible add-on devices for the installation such as a tape drive or CD-ROM drive. The original installation software may be needed if the original configuration cannot be re-created.

The plan should prepare for a total loss of the ADSM client with provisions to restore everything. Once you have provided for this worst case scenario, you will be able to recover from less serious disasters.

Bare Metal Examples

In the pages that follow examples are given for bare metal restore of AIX and Windows NT. See the bare metal restore redbook for further details or for examples of other clients. The clients explained in the redbook are OS/2 and OS/2 Warp, Windows and Windows NT, Novell NetWare, and AIX.

```
╔══════════════════════════════════════════════════════════════╗
║                      Bare Metal Restore                       ║
║                         AIX Client                            ║
║  ──────────────────────────────────────────────────────────  ║
║                                                                ║
║                           ╭─────╮                             ║
║                           │ AIX │                             ║
║                           ╰─────╯                             ║
║                                                                ║
║                      ★ ADSM Backups                           ║
║                      ★ AIX Operating Environment             ║
║                      –  Hardware Configuration               ║
║                      –  Volume Groups and Hard Disks         ║
║                      –  Communication Details                ║
║                      ★ AIX Bootable Image                    ║
║                      –  Bootable Tape Backup                 ║
║                      –  Bootable Diskette and Tape File Backup║
║                      ★ Boot Restore of AIX                   ║
║                      –  Boot System                          ║
║                      –  Restore System from Tape            ║
║                      –  Restore from ADSM                    ║
╚══════════════════════════════════════════════════════════════╝
```

Bare Metal Restore: AIX Client

To be prepared to recover an entire AIX system, collect and save the following information:

ADSM Backups

Include all data that you expect to restore with ADSM once the operating system is up and running, including the ADSM operating environment. The ADSM client options files and the Include/Exclude files will be saved as part of the system backup.

AIX Operating Environment

Collect and save the operating environment, using the DRM feature if it is available. Otherwise, keep it in hardcopy at some other location, or duplicate it to some servers where it can be readily available. The important thing is that the information exists and can be retrieved easily, not how and where it is stored.

Include the following information:

- Hardware configuration, for example, information from the lsdev and lscfg commands in AIX

- Volume groups and hard disks, for example, information from the lsvg, lslv, or lspv commands in AIX

- Communication details, for example, information from the ifconfig and netstat commands in AIX

AIX Bootable Image

A bootable image can be a bootable tape generated with the mksysb command; a set of bootable diskettes that bring up the basic system and then a full file backup; or it can be stored on a remote system and restored by the AIX network install feature.

Boot Restore of AIX

After a disaster, you must first rebuild the hardware to its original configuration.

When a tape drive is available at the client, a bootable tape is the easy way to restore. To use the network install feature, you must have the network up and running.

Restore the system backup on the re-created hardware and verify that the system is working. If it is not possible to rebuild the hardware to its original configuration, you may have to install some device drivers from an installation software copy.

After the operating system has been restored, ensure that the ADSM client options are correct and start the ADSM restore from the server. Verify the system restoration by comparing the pre- and post-disaster statistics and file information.
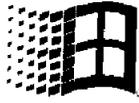
```
┌─────────────────────────────────────────────────────────────────────┐
│                        Bare Metal Restore                           │
│                        Windows NT Client                            │
│  ─────────────────────────────────────────────────────────────────  │
│                                                                     │
│                              [Windows logo]                         │
│                                                                     │
│                    ★ ADSM Backups                                   │
│                    ★ Windows NT Operating Environment               │
│                    ─  Hardware Configuration                        │
│                    ─  Communication Details                         │
│                    ─  Validation Information                        │
│                    ★ Windows NT Bootable Image                      │
│                    ─  Bootable Diskette                             │
│                    ─  Minimal NT Configuration                      │
│                    ★ Boot Restore of Windows NT                     │
│                    ─  Boot from Diskette                            │
│                    ─  Start from Recovery Partition                 │
│                    ─  Build Disk Partitions                         │
│                    ─  Restore from ADSM                             │
│                                                                     │
└─────────────────────────────────────────────────────────────────────┘
```

Bare Metal Restore: Windows NT Client

The Windows NT startup process involves two stages, a boot stage and a load
stage. The boot stage involves reading the BOOT.INI file to see which NT
partitions containing an operating system are available to be loaded. Because
the loading can be done only from a fixed disk, it is not possible to start
Windows NT entirely from diskettes. Therefore, recovering an inoperable
production Windows NT partition requires one of two methods:

1. Reinstall Windows NT from scratch, using the system installation CD, or
   possibly the emergency repair diskette that was created during the original
   installation

2. Use a diskette to initiate the startup process and have a minimal version of
   Windows NT on a different partition on the same machine. The partition
   should be on a different disk (and if possible a different disk controller) from
   the production partition. You can complete the Windows NT reload from
   there.

For more information see the redbook entitled *ADSM Server for Windows NT
Configuration and Recovery Examples* (SG24-4878).

To be prepared to recover from total loss of the operating system, collect and
save the following information:

ADSM Backups

For the ADSM backup use the default options file (DSM.OPT) of the ADSM
Windows NT client. The ADSM backup should include a copy of the registry. By
default, the ADSM Windows NT client backs up the entire registry when an
incremental backup is performed on the system partition. Use the shared static
option to avoid "fuzzy" backups.

Windows NT Operating Environment

Use the DRM to collect and save the Windows NT operating environment. If DRM is not available, keep the information in hardcopy at some other location, or duplicated to some servers where it can be readily available. The important thing is that the information exists and can be retrieved easily, not how and where it is stored.

Include the following information:

- Hardware configuration, including the partition information for the primary hard disk

- Communication details, including network specifications such as the IP address and hostname.

- Validation information, that is, information that can be used to make sure the recovery is complete, for example, a directory listing of the main directories.

Windows NT Bootable Image

We suggest using a bootable diskette and a minimal NT configuration on a recovery partition. The recovery partition should preferably be on a different physical drive and a different disk controller if possible from those of the production partition.

Install the minimal Windows NT server configuration; include TCP/IP and the ADSM Windows NT client code with this minimal configuration. The recovery system must be of the most recent Windows NT version currently in use, as the loader files may not be able to load from newer versions. Edit the BOOT.INI file and change the partition description.

Create a bootable diskette version of the Windows NT boot manager to have available if the NT system files are damaged.

Boot Restore of Windows NT

This description is based on the assumption that the Windows NT production partition has been lost and a minimum version of Windows NT is on a different partition as described in method 2 above. A complete system disaster will require a complete rebuild and restore. See the redbook entitled *ADSM Client Disaster Recovery: Bare Metal Restore* (SG24-4800).

Retrieve the information about the Windows NT client to rebuild the hardware to its original configuration. After you have recovered the hardware, recover the production platform by:

- Booting the client from the diskette boot manager

- Starting from the recovery partition prepared before the disaster, if the recovery partition is available. Otherwise go to method 1 and rebuild your system, using the system installation CD.

- Using the Windows NT disk administrator to partition the primary disk according to the information saved.

- Reformatting the primary C: drive as a FAT file system or NTFS, and then copying the system files from the boot manager diskette to the C: drive.

- With the minimal Windows NT Server system started, invoking the ADSM client and beginning the ADSM recovery.

Chapter 9. Miscellaneous Client Enhancements
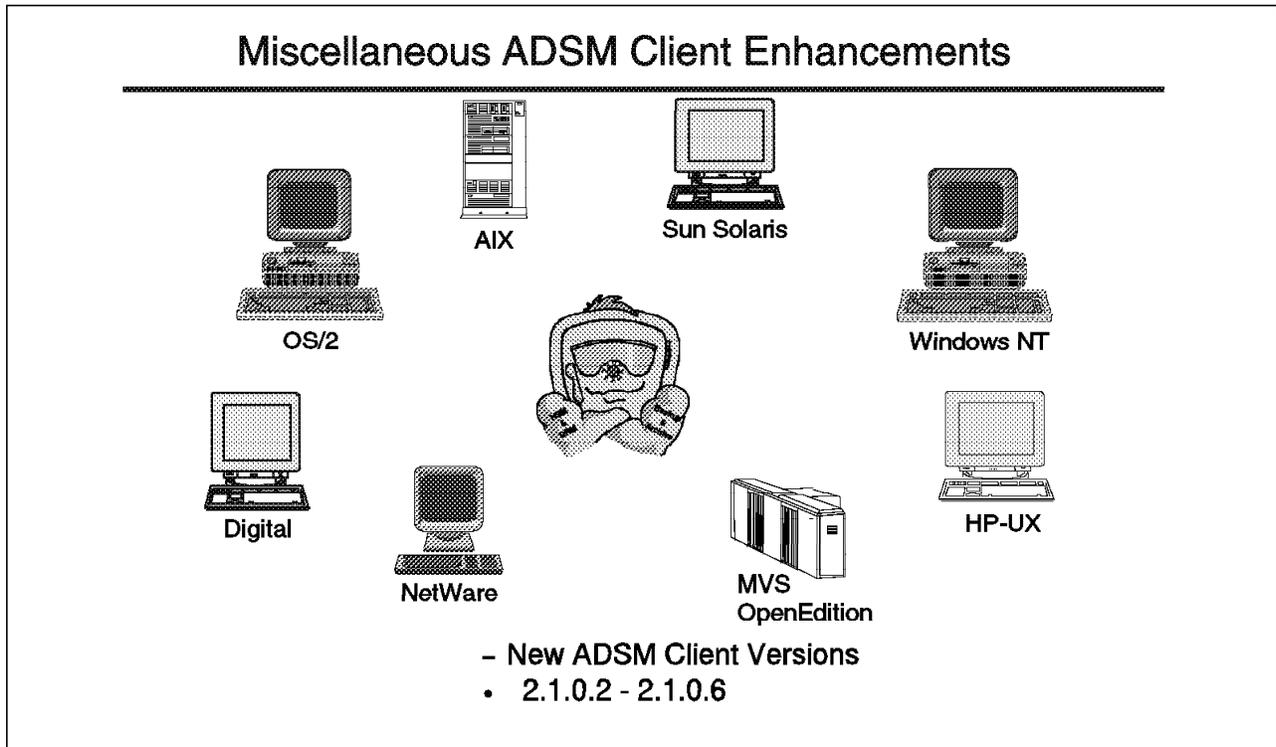
Miscellaneous ADSM Client Enhancements

In general, when an ADSM Version 2 client is made available, it is available as a client for both ADSM Version 1 and Version 2 servers. The sole exception to this is the OpenEdition MVS client, which is supported only by the Version 2 servers.

Similarly, when a new function is introduced with one of the ADSM clients, the same function will very often be available to the other clients as well. In the pages that follow, we look closer at some of the new features and functions that have been introduced through the last ADSM client updates, fix levels PTF 2 through PTF 6. Some additions have been introduced with new ADSM server functions. We identify whether a certain function is available for all clients or just for some specific clients.

We do not discuss the ADSM Version 1 clients in this book. Note, however, that once you have started using a Version 2 client on a particular node, you will no longer be able to use a Version 1 client on that node.

New ADSM Client Versions

At the end of 1996, the current ADSM client version is PTF 6, that is, ADSM Client Version 2.1.0.6. For those still running on ADSM client PTF 3 or PTF 4 versions or lower, several system fixes are available, and the clients should be updated to the newest level. As the function of the ADSM servers increases, the client code should be upgraded to take advantage of the new server functions.

With ADSM Client Version 2.1.0.6 With several important features are included, among them the support of files larger than 2 GB. The new features are described below.

Note: The PTF README files for the clients list the new features and provide advice for installation. You should always use the README information.

Year 2000

One very important fact about the ADSM backup-archive system must be mentioned here, although it is not a new client feature. With the turn of the century approaching, it is important that ADSM keep full four-digit year information and fully support the dates in the year 2000 and beyond.

Client Enhancements: General Information

Some of the information listed here can be considered an update of the ADSM documentation rather than a client enhancement. The information is important, however, and should be kept in mind.
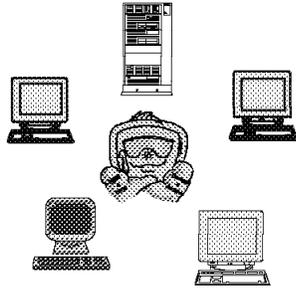
Query of Backup and Archive Output

The QUERY ARCHIVE and QUERY BACKUP commands display only the archive copies and backup versions that you own. The commands do not display files that another user owns, even if that owner has given you access to them, unless you specify -fromnode and/or -fromowner. Note that the owner of a file in ADSM storage may be different, depending on whether the file is a backup version or an archive copy. The owner of a backup version is the same as the owner of the original file. The owner of an archive copy is the person who did the archive.

Let us say, for example, that the root user backs up and archives some of your files. You would be the owner of the backup versions of your files, but the root user would be the owner of the archive copies. If you send a Q BACKUP command, you get a list of the files that the root user backed up. If you send a Q ARCHIVE command, the archive copies are not listed in the response, because the root user is the owner of those files. This is true even if the root user specifically grants you access to those files in ADSM storage.

Query with Wildcard

You can use the wildcard character ("*") as a stand-alone expression to specify all files in the current directory. When you use the wildcard character as a stand-alone expression, you must enclose it within quotation marks.

Specifying "*" in Q BACKUP will display information about backup versions for all of your files in the current directory. Specifying "*" with Q ARCHIVE will display information about all files archived from the current directory.

Restore and Retrieve of Files through the Backup-Archive GUI

When using the Restore by File Specification, Restore by Directory Tree, or Retrieve pull-down functions from the menu bar, the program's behavior must be clarified.

Once you have selected the files you want to restore or retrieve, a dialog is displayed that enables you to choose various restore and retrieve options. Under the Restore Destination section, if you choose "Restore to original location," the selected files will be restored to or retrieved from their original file system (or drive) and directories.

If you choose "Restore to following destination," all the files selected will be restored to or retrieved from the specified directory. Subdirectories are not created; all of the files are restored into one directory. This can be confusing when retrieving a set of archived directories. You must choose "Retrieve to original location" to maintain the directory structure of the retrieved files.

The Restore by Subdirectory Path function enables you to restore the entire directory tree to any location, preserving the directory structure. A future enhancement will add the capability to retrieve a directory tree from any location (preserving the directory structure).

Files Stored in File Spaces

When you back up and archive files, ADSM stores the backups and archives in a file space in ADSM storage that has the same name as the file system or virtual mount point from which the files originated.

For example, if you have a file system named /HOME, and you back up a file named DOC1 in the /HOME/JONES directory, ADSM stores the file in a file space named /HOME. However, if a root user subsequently defines /HOME/JONES as a virtual mount point, any files you back up from the /HOME/JONES directory, such as DOC2, are stored in a file space named /HOME/JONES.

If you then enter the **dsmc query backup /home/jones/*** command,, ADSM looks for files in the /HOME/JONES file space. It always looks for a file in the file space with the longest name that matches the file specification you include in a command. It finds the file named DOC2 that was backed up after the virtual mount point was defined. However, it does not find the file named DOC1 because that file was backed up before the virtual mount point was defined, and the backup was stored in the /HOME file space.

To list or restore the DOC1 file by using a command, you must specify the file space name explicitly by enclosing it in braces. For example:

- dsmc query backup {/home}/jones/*

- dsmc restore {/home}/jones/doc1

If the root user subsequently removes the /HOME/JONES virtual mount point, and you then back up additional files in the /HOME/JONES directory, the backups are once again stored in the /HOME file space. For example, if you now back up a file named DOC3 in the /HOME/JONES directory, it is stored in the /HOME file space. It is not stored in the existing /HOME/JONES file space.

However, because the /HOME/JONES file space already exists, when you try to list or restore the DOC3 file, ADSM looks in the /HOME/JONES file space for the file unless you specify the correct file space name in braces.

Note: You must specify the file space name explicitly only when there can be more than one resolution to the file specification you enter.

Cross Client Restore - Desktop Workstations

Files backed up or archived on one of the DOS, Microsoft Windows, or OS/2 clients can be restored to or retrieved from another client. However, cross-client retrieval on these platforms is subject to the following restrictions:

- When retrieving files to the DOS or Windows platforms, the files must adhere to the DOS file naming rules.

- Files stored from the OS/2 platform must use the FAT file system. If the files are using HPFS, they cannot be retrieved by DOS or Windows clients.

- A Windows client has a 256-byte limit path name.

NLM Stubs for Novell

For Novell NetWare ADSM ships special NetWare Loadable Modules (NLMs) to stub out any communication that ADSM supports but is not currently running on the system. The stubs simply export the symbols that would be exported by the communication program.

If a user starts a communication program that ADSM has stubbed out, the communication program will not load successfully. To start a communication program that ADSM has stubbed out, simply unload the appropriate stub NLM. This can occur quite innocently by loading ADSM before loading the communication NLM. The stubs provided are for TCP/IP, NetWare SAA V1.X, NetWare SAA V2.X, and PWSCS.

Workstation Identified with NODENAME

In your client system options file (dsm.sys), you can use the NODENAME option to identify your workstation to the server. If you do not include the NODENAME option in your dsm.sys file, ADSM uses the name returned by the hostname command.

You can also use the NODENAME option in your client user options file (dsm.opt) or with a command if you want to restore or retrieve files from an ADSM server that were backed up or archived for a different node. When you use the NODENAME option in dsm.opt or with a command, the node name you specify must be different from the name returned by the hostname command.

In addition, if you use the NODENAME option in your dsm.opt or with a command, ADSM prompts you for the ADSM password assigned to the node you specify (if a password is required). If you enter the correct password, you become a virtual root user and have access to all files backed up or archived from the node you specify with the NODENAME option.

Client Enhancements: Functional Enhancements

These client enhancements cover functions such as large file support, better restore statistics, and performance enhancements.

Support for Files > 2 GB
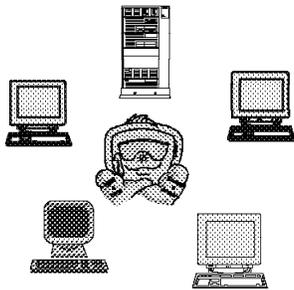
With PTF 6, the AIX 4.2, NetWare, Windows NT, and Digital UNIX clients support backup, restore, archive, and retrieve of files that are larger than 2 GB. All ADSM servers already support files larger than 2 GB.

Statistics for File Spaces > 4 GB

With PTF 6, the GUI support has been changed to allow file space statistics for file spaces larger than 4 GB. The statistics correctly report for any size files and file spaces information such as drive and file space size and number of bytes sent.

Fewer Restore and Retrieve Progress Indicators - AIX

ADSM backup-archive clients put out leader dots (...) when processing large files (backup, archive, restore, or retrieve) to show that processing is under way. Because putting out so many leader dots can affect performance on AIX, ADSM displays fewer leader dots on AIX only. Here is an example of a leader dots message:

• Normal File --> 6,348,000 /tmp/aix.tar.....................

It may appear that the file is being processed more slowly than in previous versions of the AIX client, or compared with other ADSM clients, because the dots come out more slowly on the screen. However, as ADSM now puts out only

one-fourth as many dots on the screen, the underlying command actually completes faster than in previous versions. This is especially evident with large numbers of large files.

Restore and Retrieve - Intermediate and Final Statistics

ADSM command line clients now report final statistics after a RESTORE or RETRIEVE operation.

In addition, after every 50 files have been restored or retrieved, a line is displayed that shows the total number of files that have been processed to that point.

Here is an example of the new final output:

```
Total number of objects restored:          62
Total number of bytes transferred:         4,667.6 KB
Data transfer rate:                        551.00 KB/sec
Total failures:                            0
Elapsed processing time:                   0:01:13
```

Double Buffer Reads

ADSM clients now use double buffers to read from the TCP/IP, IPX, NetBIOS, and SNA LU6.2 communication buffers. This may enhance performance on some client platforms that use these communication methods and may reduce network traffic in some cases.

Backup-Archive Client - File Modification Time (ctime)

For UNIX clients, in the Version 2 Release 1 Level 0.1 PTF, the backup and archive processes changed the ctime instead of user access time (atime). This was done because the HSM feature on AIX uses atime in calculating a file's eligibility and priority for migration. However, since the change of ctime conflicts with other existing software, from Version 2.1.0.2, UNIX backup and archive functions again perform as they did with Version 1, that is, atime is updated, but not ctime.

HSM root users can still preserve the atime (and ctime) of files on HSM active file systems by using a timestamp preserve dsmmode shell. You enter dsmmode -tp to enter the shell. Once in the shell, dsm, dsmc, and other operations will not change the time stamps.

Environment Variables

Some UNIX clients allow the use of the DSM_DIR environment variable to locate the directory where the programs are installed. Thus the system administrator can run the code from a directory other than /usr/lpp/adsm/bin. Note that dsmtca and dsm.sys must be placed in the /usr/lpp/adsm/bin directory.

The DSM_LOG environment variable locates the directory where the error log file will reside, and the DSM_CONFIG environment variable locates the directory where the config file will reside.

The following client platforms support this function: AIX, Digital UNIX, HP-UP, NCR, OS/2, Sequent, and Sinix.

Stand-alone NT Scheduler Service

The ADSM scheduler on the NT client can be run as a true stand alone service (no longer requiring SRVANY). Key features include:

- Password encryption in the registry
- Logging of key events to the NT application event log
- Random password generation for expired passwords
- Enhanced installation utility

NT Client Service Event Logging

Key events are now logged to the NT application event log. This feature is not intended to replace the ADSM scheduler or error logs. It is meant to provide administrators with a concise and distributed overview of backups scheduled through the new NT scheduler service.

Depending on whether you use detailed or brief logging, events that can be logged include:

- Password authentication failures
- Generation of new passwords when the old password expires
- Service starting and stopping
- Incremental backup success or failure
- Registry backup success or failure
- Schedule execution success or failure
- Backup and restore statistics
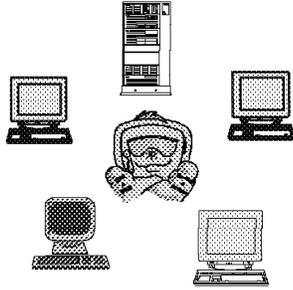
- TOTIME and FROMTIME
- Shared Memory Protocol - AIX and OS/2
- NOPROMPT with Delete Archive
- PASSWORDACCESS=generate - OS/2,
  NT, and Windows 95
- PASSWORDDIR - UNIX and OS/2

Client Enhancements: New Options

Many options have been added or changed in the last PTF updates. Some are general for all client platforms. Those that are for specific platforms are indicated.

TOTIME and FROMTIME

Two new command line options have been introduced to be used in conjunction with the todate/fromdate options to restore a file from a certain date and time of day. The TOTIME and FROMTIME options are used on the RESTORE, RETRIEVE, Q ARCHIVE, and Q BACKUP command line commands.

The FROMTIME option is used in conjunction with the FROMDATE option to specify a beginning time on the specified date. For example, with FROMTIME and FROMDATE, you can request that ADSM restore files beginning with those that were backed up on or after 6:00 a.m. on June 1, 1995.

The TOTIME option is similarly used in conjunction with the TODATE option to specify an ending time on the specified date.

You can use FROMTIME and FROMDATE with TOTIME and TODATE to request that ADSM restore, retrieve, or list only files that were backed up or archived within a particular window of time. For example, you can request that ADSM retrieve only files that were archived between 6:00 a.m. on June 1, 1995 and 11:59 p.m. on June 15, 1995.

The time must be in the format you select with the TIMEFORMAT option. When the TIMEFORMAT option is included, it must precede the FROMTIME and TOTIME options. The default for FROMTIME is 00:00:00, and for TOTIME, 23:59:59.

Note: If you want ADSM to use a value you specify for TOTIME and/or FROMTIME, you must also include TODATE and/or FROMDATE with the command. Otherwise, ADSM ignores TOTIME or FROMTIME. These options cannot be used in the options file. The AIX clients ignores the DATEFORMAT and TIMEFORMAT options and uses the local XPG4 NLS definition for these fields instead.

For backup-archive clients with a GUI only (Macintosh and Windows), sorting the active and inactive backup list and marking the first and last files to restore is a way of simulating this from-to date-time processing.

Shared Memory Protocol - AIX and OS/2

Shared Memory Protocol is for those configurations where the ADSM server and client are on the same machine (node). It is currently available with the AIX and OS/2 server. This communication protocol uses memory as the transmission vehicle for transferring data or control information between the ADSM server and the client.

It is very efficient, the fastest protocol ADSM supports, and it uses the least processor resources. Other protocols require significant processor cycles to package, transfer, and unpackage data. To use the Shared Memory Protocol, TCP/IP must be installed.

AIX

There are two options in AIX (one changed, one new) that control this communication protocol. These options, which can be specified in the dsm.sys file, are:

```
COMMMETHOD   SHAREdmem   <- Changed option
SHMPORT      1510        <- New option
```

By specifying COMMMETHOD SHAREdmem, you attempt to connect to an ADSM server running on the same machine as the client. If the ADSM server is not running or does not have Shared Memory Communications active, you will receive message ANS4038E, which indicates a TCP/IP connection failure. Shared Memory Communications uses TCP/IP to initiate the connection.

The SHMPORT option lets you specify a different TCP/IP port on which to connect to the server. The default for both the client and the server is 1510. If you are using port 1510 for some other application, however, you have to specify a different SHMPORT for both the client and the server.

OS/2

There is a new parameter to configure for use with the ADSM for OS/2 Version 2 server:

- SHMQueuename server_queue_name

Server_queue_name is the queue_name that the server listens on to establish a Shared Memory connection. The name must begin with \QUEUES and be no longer than 31 characters. It must match the server's value in the server's options file (default is dsmserv.opt). It must also be a legal OS/2 file name. The default value is \QUEUES\ADSM\DSMSERV. An example is: SHMqueuename \QUEUES\ADSM\BF.

NOPROMPT with Delete Archive

The new command line option, NOPROMPT, is only valid with the DELETE ARCHIVE command. It suppresses the ″delete archive ... ?″ confirmation prompt that appears before an archived file is deleted. Using the NOPROMPT option, you can speed up the delete operation. However, it also increases the danger of accidentally deleting an archived file you wanted to save. Use this option with caution.

For example, to suppress the confirmation prompt while deleting all files with a .tmp extension that have been archived from the c:\proj directory, type:

- Delete ARchive -nopr c:\proj*.tmp

PASSWORDACCESS=generate - OS/2, NT, and Windows 95

PASSWORDACCESS=generate keeps an encrypted password on the workstation to use during session initiation. If the password expires, a new password is automatically generated. This option is new for OS/2, NT, and Windows 95 clients.

PASSWORDDIR - UNIX and OS/2

UNIX clients can now specify the PASSWORDDIR option in dsm.sys to determine which directory ADSM should use to store the encrypted password file when the PASSWORDACCESS option is set to GENERATE. This option is also available for the OS/2 client.

This new option keeps an encrypted password on the workstation to use during session initiation. If the password expires, a new password is automatically generated.

Client Enhancements: New Options . . .
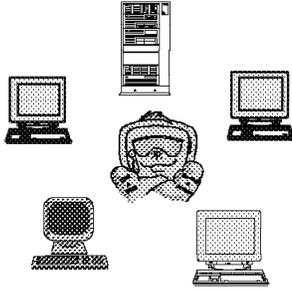
- – GUIFILESYSINFO - UNIX
- – OUTFILE
- – New Default for TAPEPROMPT
- – New Defaults for Communication Options
- – DOMAIN ALL-LOCAL - Macintosh
- – HSM - RESTOREMIGSTATE

Client Enhancements: New Options...

GUIFILESYSINFO - UNIX

This option determines whether information such as file system capacity is displayed on the initial GUI screen for all file systems (GUIF=All, which is the default) or only for local file systems (GUIF=Local). If the remote file system displayed is unreachable, ADSM must wait for the remote file system information or a timeout before displaying the initial GUI screen. This may cause a delay in the appearance of the initial GUI screen. With GUIF=Local, this delay is avoided. This option can be specified in dsm.sys or dsm.opt or on the command line when invoking the GUI, and it is relevant for the UNIX platforms.

OUTFILE

You can now specify the OUTFILE option on administrative client batchmode commands without a file specification. Specifying OUTFILE, instead of OUTFILE=filename, causes the output to be written to standard output. The output is the same format as OUTFILE=filename. This option can be useful if you are using redirection with administrative client batchmode commands.

New Default for TAPEPROMPT

The default for the TAPEPROMPT option is now NO, for performance reasons. If you want to be prompted for whether you want to wait for a tape to be mounted (if one is required for an ADSM process), you must specify TAPEPROMPT YES in dsm.opt, or you must specify TAPEP=YES as a command line option on the ADSM command.

New Defaults for Communication Options

Several communication options have new default values. Some options are only relevant for some of the clients, and some may have different defaults depending on the client. The new defaults are:

	Default Value DOS and Windows	Default Value Other Clients	Default Value Special Cases
CPICBUFFERSIZE	16 KB	31 KB	
TCPBUFFSIZE	8 KB	31 KB	Macintosh = 16 KB Windows PTF 2.1.4 = 16 KB
TCPWINDOWSIZE	16 KB	32 KB	OS/2 (DBDC lang.) = 31 KB
TXNBYTELIMIT	1024 KB	2048 KB	
IPXBUFFERSIZE	16 KB		OS/2 and Novell = 16 KB
NETBIOSBUFFERSIZE	16 KB		OS/2 = 16 KB
400COMMBUFFERBYTES	16384 bytes (16 KB)		
400SENDBUFFERBYTES	16384 bytes (16 KB)		
PROCESSORUTILIZATION			Novell = 0
SLOWINCREMENTAL			Novell = Yes

TXNBYTELIMIT controls the number of bytes, as opposed to files, that will be transferred in a group of files between transaction commit points. At the completion of transferring a file, the client commits the transaction if the number of bytes transferred during the transaction reaches or exceeds the value of TXNBYTELIMIT, regardless of the number of files transferred. The counterpart on the server is TXNGROUPMAX, which controls the number of files transferred as a group.

For DOS and FTP you may have to reevaluate the TCPWINDOWSIZE. See the relevant README file for details on how to set the option.

DOMAIN ALL-LOCAL - Macintosh

Support has been added for the DOMAIN ALL-LOCAL option. This option allows all local volumes to be selected into the domain. A local volume is regarded as a nonremovable, nonnetworked volume. Thus AppleShare, floppies, and CD-ROM volumes are not considered local volumes.

HSM - RESTOREMIGSTATE

A new RESTOREMIGSTATE option has been introduced for HSM clients AIX and Solaris 2.5. This option is specified in dsm.opt and is used only when the ADSM HSM client is installed on the workstation. It has the selections RESTOREMIGSTATE: Yes (default) or No.

Use the RESTOREMIGSTATE option to specify whether you want ADSM to restore stub files for migrated and premigrated files (placing them in a migrated state) or backup versions of migrated and premigrated files (placing them in a resident state) during a restore operation. ADSM can restore a stub file for a migrated or premigrated file only when:

- The ADSM HSM client is installed and active.
- The file exists in the migration storage pool.
- The file is backed up and migrated to the same server.

When a stub file for a migrated file or the original copy of a premigrated file is erased from a local file system, ADSM marks the migrated file for expiration the next time reconciliation is run. When the number of days specified with the MIGFILEEXPIRATION option elapses, ADSM removes the migrated file from ADSM storage.

If you set RESTOREMIGSTATE to Yes, and the migrated or premigrated file has not expired, ADSM restores the file to a stub file, regardless of whether it has been marked for expiration.

Note: The RESTOREMIGSTATE option does not support hardlinked files. If you attempt to restore a stub file for a hardlinked file, a stub file is not restored unless all of the files that are hardlinked together have been deleted from the local file system.

When one file in a set of hardlinked files is migrated, all of the hardlinked files in the set become stub files. When you issue the RESTORE command with the RESTOREMIGSTATE option, and ADSM restores a stub file for a hardlinked file, the stub file has the same name as the file that was originally migrated. Stub files are not restored for any other files that were previously in the hardlinked set of files.

With RESTOREMIGSTATE YES, ADSM restores migrated and premigrated files to stub files on the local file system during a restore operation. The files remain migrated. This is the default.

Note that a stub file created during a restore operation contains the information necessary to recall the migrated file from ADSM storage. It does not contain any leading bytes of data from the file. In addition, any recall mode previously set for the migrated file (for example, migrate-on-close or read-without-recall) is not stored in the stub file. The recall mode is set to normal for all files restored to stub files.

With RESTOREMIGSTATE NO, ADSM restores backup versions of migrated files to the local file system during a restore operation. The files become resident.

New Clients and Agents
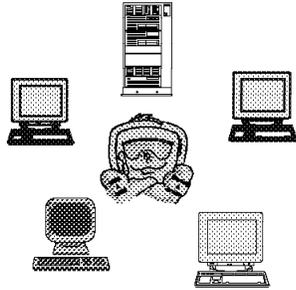
– New and Enhanced Clients
- AIX Client Support for V4.1 and V4.2
- Apple Macintosh PowerPC
- AS/400 API
- Bull DPX/2
- Digital UNIX
- GUI Interfaces Updated
- HP-UX 10.01 Client
- HSM Clients
- OpenEdition MVS Client
- Shared Memory - AIX and OS/2
- Solaris 2.5 Client
- Web Client
- Windows NT Support of API
- Windows 95
- X/Open API Compliance - AIX, HP-UX, and Solaris

– New Agents
- AFS and DFS Support on AIX
- ADSMConnect Agent for Oracle Backup
- Lotus Notes Backup Agent 4.1 Support

New Clients and Agents

New clients are sometimes introduced as part of a client PTF, and sometimes they are introduced together with a new ADSM server. The new clients described in this section include support for a new platform, a new operating system level, or new function (such as HSM).

New and Enhanced Clients

This is a list of new and enhanced ADSM clients introduced through the end of 1996.

- **AIX client support for V4.1 and V4.2**

  The AIX client has been divided into two clients, one for AIX Version 3.2, and a new one for AIX Versions 4.1 and 4.2. The clients are available through PTFs IP20800 and IP20847 for AIX V3.2 and PTFs U443912 and U444049 for AIX V4.1 and 4.2.

- **Apple Macintosh PowerPC**

  The Apple Macintosh client has been extended to support Apple Macintosh PowerPC systems.

- **AS/400 API**

  ADSM provides an AS/400 API that enables users to develop their own applications so that their AS/400 system will function as an ADSM backup-archive client.

- **Bull DPX/2**

  Bull DPX/2 has been added as a client. The Bull DPX/20 systems with AIX operating systems are already supported through the AIX client.

- **Digital UNIX**

Digital UNIX has been added as a client.

- **GUI interfaces updated**

  The Administrative GUI interfaces have been updated for OS/2, AIX, HP-UX, Windows 95, Windows NT, SunOS, and Solaris.

- **HP-UX 10.01 client**

  The HP-UX client has been divided in two, with a separate client for HP-UX Version 10.01. The HP-UX10 client also supports the X/Open API.

- **HSM clients**

  HSM, originally available only for AIX 3.2 clients, has been extended to AIX V4.1 and the new Solaris V2.5 clients. HSM on Solaris 2.5 is not supported on MP processors.

  HSM function is also provided for the Novell NetWare and Windows NT clients through the AvailHSM and OPEN/stor from Wang software. See "AvailHSM for Novell NetWare" on page 182, and "OPEN/stor for Windows NT" on page 184 for more details about AvailHSM and OPEN/stor.

- **OpenEdition MVS client**

  The OpenEdition MVS client is available for all ADSM V2.1 servers. This enhancement provides a UNIX-like client on MVS, and, with TCP/IP communication and X-window support in MVS, it also provides a GUI interface for backup-archive.

- **Solaris 2.5 client**

  The Sun Solaris client has been divided into two clients, one for Solaris Versions 2.3 and 2.4, and a new client for Solaris Version 2.5. The clients are available through PTFs IP20848 and IP20849 for Solaris V2.3 and V2.2 and PTF IP20880 for Solaris V2.5. The Solaris V2.5 client supports both the HSM feature and the X/Open API.

- **Web client**

  The Web client is made available in PTF 6 for the AIX, OS/2, Windows 95, and Windows NT clients.

  It is packaged with the backup-archive clients on these three platforms. There is a new install option (Install Web client) to select during installation of the clients. Refer to Chapter 5, "WebShell Client" on page 111 for more information.

- **Windows NT Support of API**

  The API support has been added to Windows NT.

- **Windows 95**

  The Windows NT client has been extended to include support for Windows 95, and the client has been renamed Windows 32 bit client. Some differences exist between the Windows 95 and the Windows NT clients, and they are described in "Windows 32 bit Client" on page 314.

- **X/Open API compliance, AIX, HP-UX, and Solaris**

  ADSM supports the Backup Services API specifications developed by the X/Open International Standards committee. These specifications include source procedure calls, structures, and return codes used by the AIX and Solaris clients. They answer the requirement to standardize an API between applications needing backup services and storage management vendors.

Only AIX, HP-UX 10.01 and Solaris V2.5 clients are supported.

New Agents

- **AFS and DFS support on AIX**

  In PTF 6 new AFS and DFS support (buta) is available for AIX client. It is packaged with the Open Systems Environment Support feature and is available only for AIX V4.1.4 or later. Buta provides full volume backup for the AFS and DFS, with preservation of the security attributes. This is an addition to the single file backup already available with ADSM clients.

  For further information see "AFS and DFS: Description" on page 103.

- **ADSMConnect Agent for Oracle Backup Agent**

  The ADSMConnect Agent for Oracle Backup enables ADSM to back up Oracle databases through an API into Oracle's Enterprise Backup Utility (EBU). The ADSMConnect Agent for Oracle Backup is available for the AIX clients, and it is a separate program product. It is made available in PTF 6 but has been available for testing since PTF 3.

  For more information, see "Oracle Enterprise Backup Utility - EBU (also called OBACKUP)" on page 221 or refer to *Using ADSM to Backup Databases* (SG24-4335).

- **Lotus Notes Backup Agent 4.1 support**

  The Lotus Notes Backup Agent on OS/2 now supports Lotus Notes Version 4.1 in addition to Version 3.x. The Lotus Notes Backup Agent is currently only available for OS/2. It requires the ADSM OS/2 API which runs on the OS/2 Warp 3.0 or Warp Server and provides incremental backup and recovery of Lotus Notes databases. For further details see "ADSM and Lotus Notes" on page 240.

Chapter 10. Miscellaneous Server Enhancements

Miscellaneous ADSM Server Enhancements

This chapter describes the enhancements for the ADSM Version 1.x and the ADSM Version 2.x servers (x = different releases and service levels).

ADSM Version 1.x Server

The enhancements for the ADSM Version 1 server can be implemented with the appropriate PTF level.

The implementation steps are:

1. Establish the release- and service-level of your installed ADSM server version.
2. Contact your local IBM customer service representative to order the desired PTF.
3. Upgrade your ADSM server to the appropriate PTF level.

ADSM Version 2.x Server

To implement the enhancements for the ADSM Version 2 server, follow the steps for implementing the ADSM Version 1.x server.

Now is a good time to upgrade your server to the current available service level to take advantage of all enhancements!

157

ADSM Version 1.x Server Enhancements
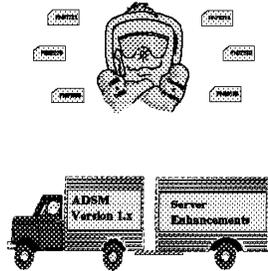
Enhancement	ADSM Version 1.x Server							
	MVS	VM	VSE	AIX	OS/2	OS/400	Sun	HP-UX
DB Salvage Utility	✳	✳	✳	✳	✳	✳	✳	✳
QUERY EVENT		✳						
EXPINTERVAL		✳						
MESSAGEFORMAT	✳	✳						
Multiple Migration Processes	✳	✳	✳	✳	✳	✳	✳	✳
RENAME FILESPACE	✳	✳	✳	✳	✳	✳	✳	✳
AUDITDB FILE=...	✳	✳						
DISMOUNTING MSG	✳	✳						
TXNGROUPMAX	✳	✳	✳	✳	✳	✳	✳	✳
STGPOOL THRESHOLD	✳	✳	✳	✳	✳	✳	✳	✳
EXPORT Performance	✳	✳	✳	✳	✳	✳	✳	✳
STATUSMSGCNT	✳	✳	✳	✳	✳	✳	✳	✳
DSN PREFIX with Periods		✳						
STGPOOL Limit	✳	✳						
EXPIRE INVENTORY QUIET=Y/N	✳	✳	✳	✳	✳	✳	✳	✳

ADSM Version 1.x Server Enhancements

This graphic lists all of the ADSM Version 1.x server enhancements and the
ADSM server platform on which they are available.

DB Salvage Utility

PTF 08
APARs PN49454, PN57592, PN57505, PN48720, and others

The ADSM server program does not provide all facilities needed to easily
recover the server database should recovery be necessary. Rapid recovery
currently requires that database and recovery log volumes be copied with
external utilities while the server is not running. In addition, the recovery
process may require long database and data storage volume audits.

With this PTF, the server has been enhanced to provide functions that facilitate
planning for and executing ADSM server database recovery. Should you find
that you have to recover server database contents because of a device or
system failure, these enhancements will assist you in successfully restoring your
server to proper operation.

For detailed information, refer to the appropriate APAR.

QUERY EVENT

PTF 08
APAR PN51721 and others

The QUERY EVENT command has been enhanced so that the status field always contains a value. In addition to the values documented in the *ADSM Administrator's Reference* for your server platform, the following new values can be used:

- FUTURE - specifies that the beginning of the startup window for the event is in the future

- PENDING - specifies that the QUERY EVENT command was issued during the startup window for the event, but execution of the scheduled operation has not yet begun

- UNCERTAIN - specifies that the state of the event cannot be determined (for example, if the event record was deleted)

EXPINTERVAL

PTF 08
APAR PN51821 and others

The ADSM server executes inventory expiration automatically every hour. The expiration function requires significant resources when executing on large ADSM servers. With this PTF, EXPINTERVAL is available as a server option in the server options file.

You can specify "EXPINTERVAL hours," where hours indicates the number of hours between automatic inventory expiration runs. The minimum value permitted is 0; the maximum value is 336. With a value of 0, you have to start the inventory run manually with the EXPIRE INVENTORY command.

MESSAGEFORMAT

PTF 08
APAR PN52270 and others

Currently only the first line of a multiline message gets the message header. Because message suppression is possible on MVS only when a message header is present, only the first line of a multiline message will be suppressed.

Function has been added to put message headers in all lines of a multiline message. This function is controlled by a new option, MESSAGEFORMAT, which is specified in the server options file. Option 1 (the default) indicates that only the first line of a multiline message gets the message header. Option 2 places a message header in all lines of a multiline message.

Multiple Migration Processes

PTF 08
APAR PN54180 and others

Migration of files from one ADSM storage pool to another is currently performed through a single process. One consequence of single-process migration is that I/O to the target pool is restricted to a single output device, which can limit the migration rate. For example, if migration is performed from a random access

storage pool to a tape storage pool, the migration rate may be limited by the speed of the tape drive used for writing to the target pool. Because of this restriction, migration from a storage pool may not be able to keep up with the influx of files into that storage pool; that is, the storage pool is filled faster than it can be emptied.

To address this problem, the operation of migrating files from random access storage pools has been enhanced. A new attribute has been added that indicates the number of processes which should be used for migration. During migration, these processes are performed in parallel, thus providing the potential for improved migration rates.

As an example, suppose that you wanted to improve the migration rate between two storage pools. Also suppose that four tape drives are available for migration processing, one random access storage pool, and one sequential access storage pool. BACKUPPOOL is the random access storage pool, and TAPEPOOL is the sequential access storage pool assigned to a tape device class. Files located in BACKUPPOOL are migrated to TAPEPOOL. To improve the migration rate from BACKUPPOOL to TAPEPOOL, you might set the number of migration processes for BACKUPPOOL to 4 so that all four tape drives are used.

The amount of parallelism that can be achieved by multiple migration processes may be limited by the number of nodes with migratable data. When selecting the next node for which to migrate files, a process will skip any nodes for which migration is already in progress. Because each process migrates files for a different node, a process will not be able to perform migration if nodes with migratable files are not available.

When determining the number of migration processes to use for migration to sequential access storage pools, consider the number of available drives. Each migration process will use a drive, so the number of migration processes should not exceed the number of drives that can be dedicated to migration processing into the target pool.

The DEFINE STGPOOL and UPDATE STGPOOL administrative commands have a new parameter, MIGPROCESS=numprocesses (numprocesses is an integer in the 1 - 999 range). For the DEFINE STGPOOL command, the default value for MIGPROCESS is 1.

RENAME FILESPACE

PTF 08
APAR PN57392 and others

When a filespace is imported into the ADSM server, the import process generates a new name for the filespace if an existing filespace with the same name already exists for the client node in the server database.

The name generated for the imported filespace may not be appropriate for the client node, or it may conflict with an existing filespace on the client that is not yet known to the ADSM server. The filespace has not yet been backed up to the server. You can use the new RENAME FILESPACE command to rename a client filespace that is not appropriate for the client or conflicts with client filespace names that are unknown to the server.

Here is the command syntax:

REName FIlespace nodename filespacename newfilespacename

AUDITDB FILE=...

PTF 09
APAR PN53459 and others

The DSMSERV AUDITDB process can issue a large number of messages. On MVS systems these messages can fill up the space for the job and cause system problems.

The ADSM server's AUDITDB processing has been modified to allow a new command parameter, FILE, where you specify a valid file name for the system on which the server is running. The file will be dynamically created if it does not exist and will be overwritten if it exists. The AUDITDB messages are written to this file. If the redirection fails, the messages will be issued back to the console.

On VM servers, the format for the file name is FILE=fn.ft.fm, where fn = file name; ft = file type; and fm = file mode for CMS minidisk files.

Dismounting Message

PTF 10
APAR PN61585 and others

The ADSM VM and MVS servers do not issue messages about volume update, when tape volumes are dismounted. Customers must know the update status of volumes because they cannot easily automate operations to be performed on tapes that are updated by the ADSM server.

With this PTF each time a sequential storage pool volume is dismounted, the ADSM VM and MVS servers issue a message indicating whether or not they updated the volume while is was mounted.

TXNGROUPMAX

PTF 14
APAR PN67795 and others

Specifying the TXNGROUPMAX in the server options file controls the number of files that will be transferred as a group between a client and a server between transaction commit points. You may be able to improve the performance of the BACKUP, ARCHIVE, RESTORE, AND RETRIEVE commands by using a high value for TXNGROUPMAX.

The files transferred for the BACKUP, ARCHIVE, RESTORE, AND RETRIEVE commands are either actual files or directories or both. Each actual file and each directory are counted as one file.

To obtain a performance improvement on RESTORE and RETRIEVE commands, a client PTF is also needed.

TXNBYTELIMIT controls the number of bytes, as opposed to files, that will be transferred in a group of files between transaction commit points. On completion of file transfer, the client commits the transaction if the number of

bytes transferred during the transaction reaches or exceeds the value of TXNBYTELIMIT, regardless of the number of files transferred.

The TXNBYTELIMIT parameter is implemented in the following ADSM Version 1 client PTFs:

- IP20258   AIX/6000
- IP20261   DOS
- IP20262   OS/2 (2.0)
- IP20263   OS/2 (1.3)
- IP20264   Windows
- IP20265   Novell
- IP20266   HP
- IP20267   DEC
- IP20268   SCO
- IP20269   Sun (disk)
- IP20270   Sun (tape)

Storage Pool Threshold

PTF 15
APAR PN72319 and others

Because the high migration threshold for a storage pool cannot be set to 0, it is impossible to use migration to empty a storage pool when it only has a few files in it and the %Migr value rounds to 0.

After installation of this PTF, the high migration threshold on DEFINE STGPOOL and UPDATE STGPOOL can now be a value from 0 to 100. Migration can then completely empty a storage pool.

Export Performance

PTF 16
APAR PN73196 and others

The ADSM server EXPORT function has been modified to improve performance.

STATUSMSGCNT

PTF 16
APAR PN77273 and others

Customers who have very large databases can see too many status messages when dumping or loading their databases. A new server option has been added to enable customers to control how often these messages are output:

STATUSMSGCNT xxx (xxx = a value between 1 and 10000; default 10)

The value specified is multiplied by 1000, and the result is used to determine how often (after how many database records are processed) status messages will be output.

DSN PREFIX with Periods

PTF 16

APAR PN77286 and others

Periods are now allowed in the PREFIX value. For device class cartridges, if a value is specified for PREFIX, it must meet the following conditions: The value, which can contain up to eight characters (including periods) must be made up of qualifiers. For example, the following value would be acceptable: AB.CD2.E.

Storage Pool Limit

PTF 17

APARs PN80136, PN73847, and others

PTF 17 increases the storage pool limit from 250 to 999.

Expire Inventory QUIET=Y/N

PTF 17

APAR PN81092 and others

You can perform expiration processing with the new QUIET option. The option enables you to suppress a set of detailed informational messages that would normally be issued when you make changes to policy information such as deleting a management class, deleting a backup copy group, or deleting an archive copy group. Instead of these detailed informational messages, a summary message is issued only when files have been deleted and either the default management class or the retention grace period for the domain has been used to expire files.

When expiration processing is started automatically with the EXPINTERVAL option in the server options file, use the new EXPQUIET YES or NO option to select the information level you need. YES = only minimal summary messages. NO = detailed messages (the default).

The EXPIRE INVENTORY command now has a new parameter, QUIET=YES or NO, to manage this function when the expiration process is started manually.

ADSM Version 2.x Server Enhancements
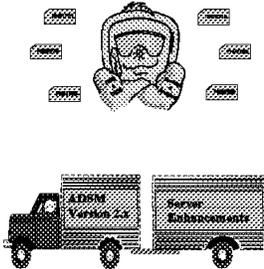
Enhancement	ADSM Version 2.x Server							
	MVS	VM	AIX	OS/2	OS/400	NT		
EXPORT Performance	✱	✱	✱	✱	✱			
DSN Prefix with Periods	✱	✱	✱		✱			
STATUSMSGCNT	✱	✱	✱	✱	✱			
DEFINE DEVCLASS New Options	✱							
NOAUDITSTORAGE	✱	✱	✱	✱	✱			
DSMSERV INSTALL PARM Field	✱							
STGPOOL Limit	✱		✱					
TAPE Dismount Delay	✱							
DELETE VOLUME Performance	✱		✱					
EXPIRE INVENTORY QUIET=Y/N	✱		✱					
COLLOCATION by FILESPACE	✱	✱	✱	✱	✱			
STGPOOL Reclamation Mountlimit=1	✱	✱	✱	✱	✱			
RESTORE VOL/STGPOOL Performance	✱	✱	✱	✱	✱			
BARCODE Support SCSI Tape Library			✱	✱				
Shared Memory Protocol			✱	✱				
PTHREADS			✱					
IBM 3570 Support					✱			
CLIO/S	✱							

ADSM Version 2.x Server Enhancements

This graphic lists all of the ADSM Version 2.x server enhancements and the ADSM server platform on which they are available.

Export Performance

PTF 02
APAR PN73196 and others

The ADSM server export function has been modified to improve performance.

DSN PREFIX with Periods

PTF 03
APAR PN75677 and others

Periods are now allowed in the PREFIX value. For device class cartridges, if a value is specified for PREFIX, it must meet the following condition: The value, which can contain up to eight characters (including periods), must be made up of qualifiers. For example, the following value would be acceptable: AB.CD2.E.

STATUSMSGCNT

PTF 03
APARs PN76442, IC11464, and others

Customers who have very large databases can sometimes see too many status messages when dumping or loading their databases. A new server option enables customers to control how often these messages are output:

STATUSMSGCNT xxx (xxx = a value between 1 and 10000; default 10)

The value specified is multiplied by 1000, and the result is used to determine how often (after how many database records are processed) status messages will be output.

DEFINE DEVCLASS New Options

PTF 07
APAR PN73862 and others

ADSM invokes dynamic allocation to allocate sequential data sets used as volumes. Exactly where the volumes are placed depends on the system configuration of an installation.

Two options, UNIT and VOLSER, have been added for device type=FILE device classes to enable ADSM to allocate FILE volumes on specific MVS devices.

NOAUDITSTORAGE

PTF 07
APARs PN77064, IC12429/IX57138, and others

When an audit license runs on an ADSM server, it calculates the amount of storage in use on the server. For very large servers managing very large amounts of data, this calculation can take a great deal of real time and a significant amount of processor time and can stall all other server activity.

An option has been added to the server that enables you to indicate whether or not you want the audit license process to calculate storage use. The option is NOAUDITSTORAGE, which is to be placed in the server options file, followed by a restart of the server.

DSMSERV INSTALL PARM Field

PTF 07
APAR PN78840 and others

Fully qualified data set names may be long. The PARM field is limited to 100 characters. Customers would like to provide files containing the candidate recovery log and database volume names when installing the server. The DSMSERV INSTALL has been changed with this PTF to reflect the PARM field.

For the new syntax please see the APAR description.

Starting with ADSM Version 2, the server does not format VSAM linear data sets as part of the INSTALL or DEFINE commands. The data sets must be allocated and formatted outside the server. ASAMPLIB members are provided to help!

Storage Pool Limit

PTF 07
APAR PN80135, IC12255/IX57134, and others

PTF 07 increases the storage pool limit from 250 to 999.

Tape Dismount Delay

> PTF 07
> APAR PN81941 and others

Users of ADSM tapes could experience a delay in the dismount of tapes when multiple tapes are being dismounted. The CLOSE and DEALLOCATE routines are changed to reduce the possibility of enqueue delays.

DELETE VOLUME Performance

> PTF 09
> APAR PN82683 and others

The processing for DELETE VOLUME has been enhanced to group the delete transactions into larger batches for improved throughput for multiple concurrent DELETE VOLUME DISCARDDATA=YES operations.

Expire Inventory Quiet=Y/N

> PTF 09
> APAR PN83440 and others

Expiration processing can be performed with the new QUIET option. This option enables you to suppress a set of detailed informational messages that would be issued when you make changes to policy information such as deleting a management class, deleting a backup copy group, or deleting an archive copy group. Instead of detailed informational messages, a summary message is issued only when files have been deleted and either the default management class or the retention grace period for the domain has been used to expire files.

When expiration processing is started automatically with the EXPINTERVAL option in the server options file, use the new EXPQUIET YES or NO option to select the information level you need. YES = only minimal summary messages. NO = detailed messages (the default).

The EXPIRE INVENTORY command now has a new parameter, QUIET=YES or NO, to manage inventory expiration when the expiration process is started manually.

COLLOCATION by FILESPACE

> PTF 10
> APARs PN83701, IC14838/IX60878 and others

Customers with clients that have a large number of file systems require collocation at the filespace level to minimize tape mounts.

This PTF allows further granularity for collocation. If a primary sequential access or copy storage pool is defined appropriately, the server attempts to put data for one node and filespace on one volume. If a node has multiple filespaces, the server attempts to place data for each filespace on different sequential volumes in the storage pool.

To enable collocation by filespace, specify the new FILESPACE value for the COLLOCATE parameter on the following commands:

- DEFINE STGPOOL ... COLLOCATE = NO | YES | FILESPACE

- UPDATE STGPOOL ... COLLOCATE = NO | YES | FILESPACE

Storage Pool Reclamation Mountlimit=1

PTF 10
APARs PN88447, IC14216/IX60866, and others

Offsite reclamation should work with a mountlimit of 1 because the reclamation process copies the files from other volumes in other storage pools. Reclamation can now be performed for an offsite copy storage pool volume that resides in a storage pool associated with a devclass with a mountlimit of 1.

RESTORE VOLUME/STGPOOL Performance

PTF 10
APAR IC14818/IX60827, and others

Use the RESTORE VOLUME and RESTORE STGPOOL commands to restore files from a copy storage pool to a primary storage pool. If the copy storage pool is noncollocated and the primary storage pool is collocated, restore processing can be very slow.

When restoring to a collocated primary storage pool, the server restores files by node and filespace. Thus collocation is maintained in the target storage pool and tape mounts are minimized during restore processing. If the copy storage pool is noncollocated, grouping of files by node and filespace is time consuming and leads to degraded performance.

RECOMMENDATION:

A code change has been made that provides much better performance for restoring files from a noncollocated copy storage pool to a collocated primary storage pool. Even with this performance enhancement, however, restore processing to a collocated pool is considerably slower than restoring to a noncollocated pool. To achieve the best performance when restoring files to a collocated pool, customers should consider restoring to a random-access disk storage pool and then allowing the files to migrate to the collocated primary pool. By restoring to the disk pool, restore performance can be greatly improved, because files need not be grouped as they are restored. Furthermore, migration from a disk pool to a collocated pool is much more efficient than performing a storage pool restore to a collocated pool. Thus, restoring files to a disk pool and then allowing the files to migrate to a collocated pool is usually faster than restoring directly to the collocated pool. To reduce the requirements for disk space, it is also possible for the restore and migration operations to run concurrently.

Barcode Support SCSI Tape Library

This support is currently available for ADSM for AIX Version 2, Release 1, PTF level 1 or higher. It is also included in the ADSM for OS/2 Version 2 server base product.

The barcode support exploits barcode hardware to make ADSM server operations more efficient; less human intervention is required for installing ADSM with larger SCSI libraries.

With the addition of barcode support, ADSM server operations and installations are enhanced because:

- Cartridges in a SCSI library can be labeled by a single command, with no requirement to enter volume names from the console.

- CHECKIN processing for a SCSI library can be completed in seconds, even when the number of cartridges is large.

- AUDIT LIBRARY processing can also be performed in seconds.

- CHECKOUT processing is faster and will also notice conditions that might indicate that an audit is required.

New parameters have been added to the DSMLABEL, CHECKIN, AUDIT, and CHECKOUT commands to support this enhancement.

For detailed information please contact your local IBM software service center.

Shared Memory Protocol

The Shared Memory Protocol is for those configurations where the ADSM server and client are on the same machine (node). The protocol is currently available with the AIX and OS/2 servers. It uses memory as the transmission vehicle for transferring data or control information between the ADSM server and the client.

The Shared Memory Protocol is the fastest protocol that ADSM supports, and it uses the least processor resources. Other protocols require significant processor cycles to package, transfer, and unpackage data. To use the shared memory protocol, TCP/IP must be installed.

PTHREADS

PTHREADS is an enhancement to the ADSM Version 2 server. The enhancement provides improved support for shared multiprocessor (SMP) machines running AIX 4.1.4 or later and eliminates the ADSM kernel extension, making configuration easier. Use of PTHREADS makes ADSM a little more efficient because you use threads for much of the ADSM server's multiprocessing, rather than UNIX processes (UNIX processes have more overhead!).

PTHREAD support is available on ADSM for AIX Version 2.1.5.6 under AIX Version 4.1.4 or later. ADSM for AIX Version 2.1.0 does not support PTHREADS.

IBM 3570 Support

APARs SA56721 and SA57127

If the ADSM server is to use an IBM 3570 Cartridge Tape Device, there is a special consideration that you must address. The 3570 device contains a feature that permits an application to move to a specific file on the tape volume. When a file is created, its position on tape can be recorded. During subsequent read operations, the application can request that the 3570 device "fast forward" to that position instead of reading each file on the tape. This fast forwarding reduces the total time required to restore data from tape. Within a given AS/400 system, only one application can provide an exit program that will be given control at the AS/400 tape management exit point.

If Backup Recovery and Media Services/400 (BRMS/400) is using the exit point, the location of a file on tape is kept in the BRMS media information. We

recommend that you define a USRDFN library to permit BRMS/400 to manage the 3570 for the ADSM server. ADSM will gain any performance benefits arising from the ability of BRMS/400 to fast forward to a file.

If the AS/400 tape management exit point is available, the ADSM server attempts to use it. The server keeps the location of a file on tape in its database, thus enabling ADSM to gain any performance benefits arising from its own capability to "fast forward" to a file.

If the Media and Storage Extensions (MSE) feature has not been installed on your system, you will still be able to use the 3570. You will not, however, be able to use the fast forward capabilities of the device.

CLIO/S

PTF 12

ADSM Client Input/Output Sockets (CLIO/S) support uses the CLIO/S fast transfer feature to provide high-speed connectivity between AIX client and MVS server environments. ADSM CLIO/S support may increase performance between an AIX client and MVS server and can also reduce processor utilization of the TCP/IP address space on MVS.

ADSM CLIO/S support is provided through the new MVS server option, CLIOPort, and client option COMMmethod CLIO. For details on how to use these options, refer to the README file.

These are the prerequisites for ADSM CLIO/S support:

- ESCON or parallel-channel attachment between the client and server platform
- CLIO/S Version 2.2
- TCP/IP Version 3.1 or 3.2 (Version 3.2 provides reduced processor utilization)

  Interlink TCPaccess is not supported.
- ADSM Version 2 server with PTF level 12 (December '96)
- AIX Version 2 client with PTF level 5

For more information about CLIO/S, refer to the *CLIO/S V2.2 User's Guide and Reference* (GC28-1676)

Chapter 11. ADSM and Complementary Products

ADSM and Complementary Products

In this chapter we describe some vendor and IBM products that work in conjunction with ADSM. Some products, such as Tandem Guardian and OPEN/stor, use the API to use ADSM as a data respository, such as Tandem Guardian and OPEN/stor. We also include two service offerings for CRAY and Fujitsu clients. Some products provide functions that ADSM may take advantage of, such as remote copy and Computer Network Technology (CNT) Channelink.

ADSM CRAY UNICOS Client
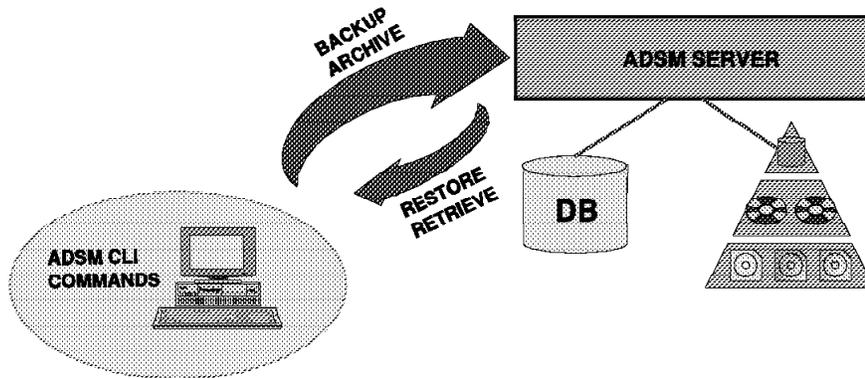
ADSM CRAY UNICOS Client

The ADSM CRAY UNICOS client is available as a service offering from Storage Manufacturing Information Technology (SM IT) Mainz in Germany. It is not a part of the program product for ADSM.

The service offering provides a backup-archive client, an administrative client, and maintenance and support.

CRAY UNICOS 8 or 9

The CRAY UNICOS client requires a UNICOS 8 or 9 operating system running on CRAY UNICOS SMP models J90, C90, YMP, or T90.

The CRAY UNICOS client is based on the ADSM Version 2.1 client and works like the other ADSM UNIX clients. It requires TCP/IP as the communication protocol.

Maintenance and Support

The service offering also includes the following maintenance and support:

- Maintenance and trouble shooting are provided for the first six months. Maintenance thereafter can be provided at an additional cost.

- A support telephone line and mailbox are provided for addressing your problems directly to SM IT Mainz.

ADSM CLI

CRAY UNICOS has a CLI for the backup and archive clients. You can back up and restore as well as archive and retrieve files. You use the same commands as you would for any other ADSM CLI.

Future Administrative Client with CLI

SM IT Mainz intends to support the administrative client with the command line interface.

Order Information

The code is shipped as object code only. Contact SM IT Mainz for ordering.

More Information

If you need more information, contact SM IT Mainz in Germany.

ADSM Fujitsu UXP/V Client

The ADSM Fujitsu UXP/V client is available as a service offering from Storage Manufacturing Information Technology (SM IT) Mainz in Germany. It is not a part of the program product for ADSM.

The service offering provides both features and maintenance.
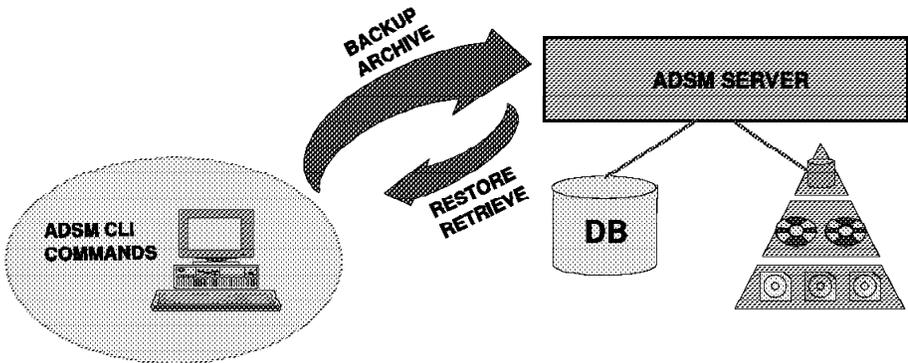
UXV/P on VPP300 or UXV/M on VPS600

The Fujitsu UXP/V client requires the UXV/P operating system running on a Fujitsu VPP300 series model or the UXP/M operating system running on a Fujitsu VPS600 series model. It uses TCP/IP as the communication protocol.

Maintenance and Support

The service offering also includes the following maintenance and support:

- Maintenance and trouble shooting are provided for the first six months. Maintenance thereafter can be provided at an additional cost.

- A support telephone line and mailbox are provided for addressing your problems directly to SM IT Mainz.

ADSM CLI

The client is based on the ADSM Version 2.1 client and works like the other ADSM UNIX clients.

Fujitsu UXP/V has a CLI for the backup-archive and administrative clients.

API

This service offering also includes an API.

Order Information

The code is shipped as object code only. Contact SM IT Mainz for ordering.

More Information

If you need more information, contact SM IT Mainz.
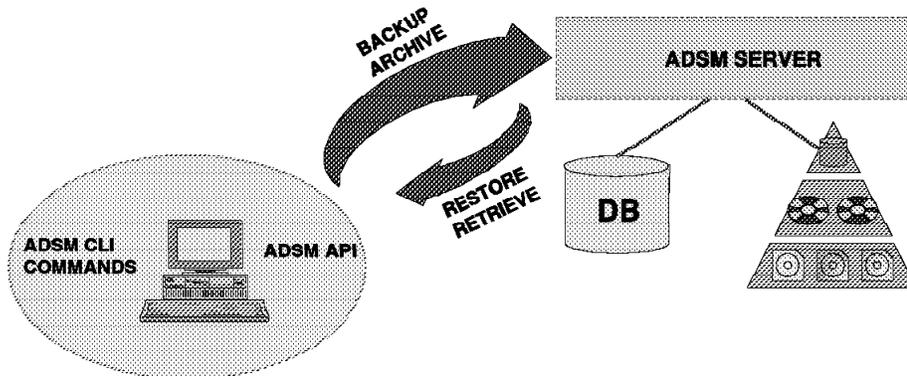
OpenVMS Client for ADSM

The Storage Solutions Specialist Inc. (SSSI) Archive Backup Client enables Digital OpenVMS users to back up, archive, restore, and retrieve their data through the ADSM API. This backup-archive client is called the OpenVMS client.

An OpenVMS client can use ADSM services through the ADSM API. You can back up, archive, restore, and retrieve files stored on the ADSM server as a logical extension to the online OpenVMS ODS-2 file system. You do not have to be concerned about the tape media or storage libraries; ADSM manages them for you.

OpenVMS Version 5 and 6

The OpenVMS client requires OpenVMS V5.5-2 or higher for VAX systems and OpenVMS V6.1 or higher for Alpha systems.

You need an ADSM server that uses TCP/IP as the communication protocol and supports the ADSM Client API Version 2, Release 1, Level 2 or higher.

Command Interface

With a few exceptions, the OpenMVS client gives you the same benefits and features as the other backup-archive clients.

Though the command line interface, you can tell ADSM which operation you request. You use the BACKUP or ARCHIVE command to store data in the ADSM storage pools, and you use RESTORE or RETRIEVE command to get a file back from the storage pool. With the SHOW command you can display information from the ADSM database, such as archive files, backed up files, management classes, and filespaces. The OpenVMS client does not support the GUI.

You must have a password to log on to the ADSM server. You can restore or retrieve only those files that you previously backed up or archived. The OpenVMS client only supports the GENERATE passwordaccess option.

You can schedule your backup and archive operations in two ways:

- Use the OpenMVS scheduling mechanisms, either batch queues or the POLYCENTER scheduler

  or

- Use third-party scheduling products.

ADSM central scheduling is not supported, but SSSI intends to support it in a maintenance upgrade.

Recovering a system disk is often a complex and time-consuming operation. If you have used ADSM to successfully back up your files, you can produce a bootable system disk from another system. Once you have a bootable system you can restore all of your OpenVMS client files.

Order Information

You can order the OpenVMS Client for ADSM V1.0 from SSSI, or you can order it from IBM (product number is 5620-ART).

More Information

For more information, see the Archive Backup Client home page at http://www.storsol.com.
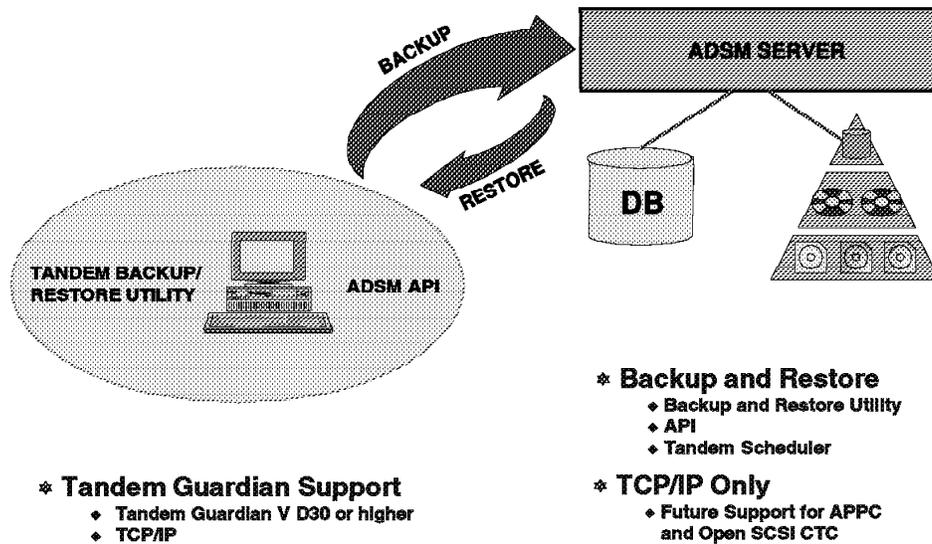
Tandem Guardian Client

Experts en Traitement de l'Information (ETI), Montreal, Canada, offers you an ADSM client solution for the Tandem Guardian platform. The Tandem Guardian Client offering is based on the ADSM API ported to the Tandem Guardian platform and is used in conjunction with ETI's BackHome! for ADSM product.

Tandem Guardian Support

You can run the Tandem Guardian client by using Backhome! on any hardware that supports Tandem Guardian Version D30 or higher. You can use any ADSM server that supports TCP/IP as the communication protocol.

Backup and Restore

By using Backhome! for ADSM you can back up and restore any Tandem file or disk to ADSM storage pools that the Tandem Backup/Restore Utility can back up and restore. You cannot archive or retrieve files.

BackHome! uses requests to define a backup or restore function. Each request contains:

- Specification of which files or volumes you want to back up

- ADSM server information that tells the Tandem Backup/Restore Utility which ADSM server to use to store the data

- Specification of which function to use: backup or restore

Once the request is complete, it is stored in a request file. You can immediately run any backup or restore request that has been defined in the request file, or you can queue a request to the Tandem Scheduler.

You use the Tandem Scheduler to request automated backups. ADSM central scheduling is not supported.

TCP/IP Only

Although only TCP/IP is supported today, ETI intends to provide support for APPC and Open SCSI CTC communication protocols.

Order Information

Backhome! for ADSM is available as a separate product or as a feature of the BCOM-FS 6.2 product. Please contact ETI for ordering.

More Information

For more information about the BackHome! for ADSM client, see the ETI home page at http://www.etiusa.com.
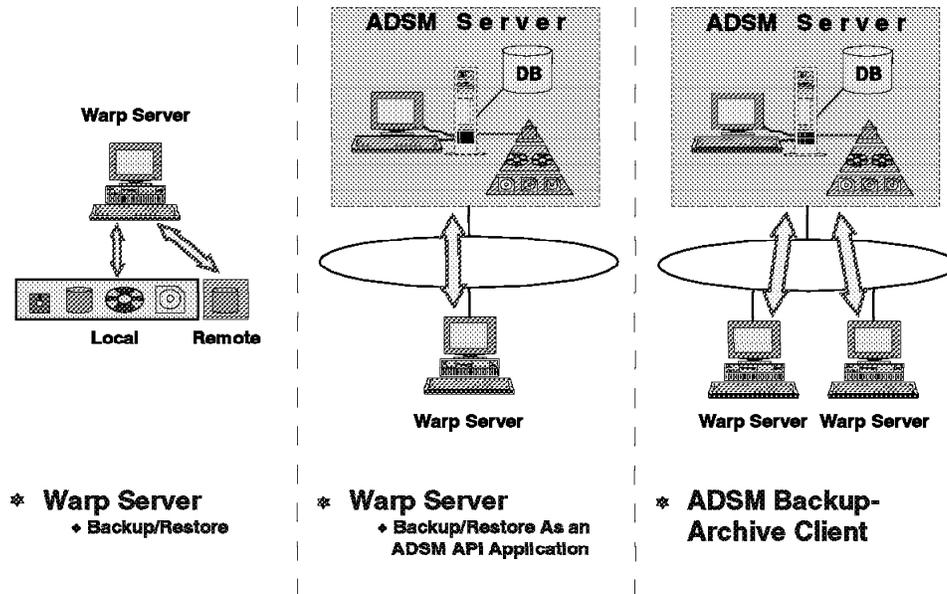
OS/2 Warp Server Backup/Restore and ADSM

OS/2 Warp server data can be distributed across multiple drives and drive types and may have system or security information associated with it. The data may have different backup and recovery requirements. You can extend the backup and restore feature that comes with the OS2 Warp Server by using ADSM.

Warp Server Backup/Restore

The OS/2 Warp Server Backup/Restore utility manages backups for data on:

- Diskette drives
- Locally attached hard drives
- LAN alias hard drives
- Remotely attached hard drives
- Locally attached manual change tape drives
- Locally attached manual change optical drives

The utility provides you with default backup and restore routines and a GUI for quick and easy setup of backup, restore, and recovery procedures.

Warp Server Backup/Restore As an ADSM API Application

For more powerful backup and restore features, use the OS/2 Warp Server Backup/Restore utility as an ADSM API application. The utility is responsible for the backup and recovery of your data, but you use ADSM storage pools to store your data. As soon as the data is associated with the ADSM storage media, the Backup/Restore utility releases control of the data, and the ADSM server is responsible for managing and protecting the data. ADSM can move the data between its storage volumes according to the storage hierarchy and migration thresholds specified.

ADSM Backup-Archive Client

In an environment where you have a small number of servers and a small amount of data to manage, the Backup/Restore utility is ideal. When your environment grows, and the number of clients and/or the amount of data increases, you may want to consider using the standard ADSM OS/2 client to obtain all of the advantages of a centrally managed ADSM environment.

More Information

For more information about the OS/2 Warp Server Backup/Restore utility, see *Using ADSM to Back Up OS/2 LAN Server and Warp Server* (SG24-4682).
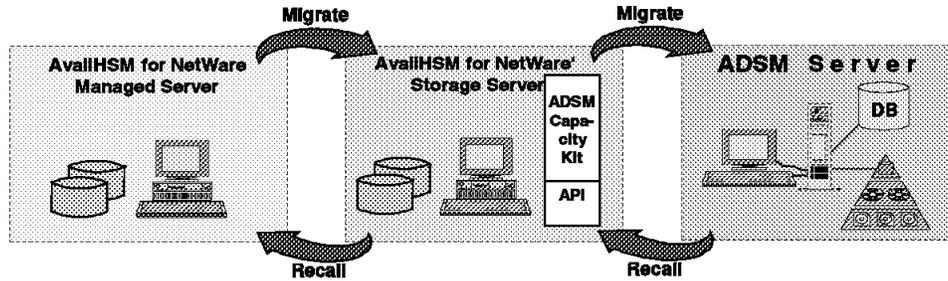
AvailHSM for Novell NetWare

The ADSM Novell NetWare client supports the backup-archive client only. ADSM does not provide an HSM solution for the NetWare environment. The ADSM space management solution for NetWare servers is the AvailHSM product. AvailHSM is available as a separate product from Wang Software Storage Management Group, Inc. This product complements the backup and archive functions provided by the ADSM NetWare client. AvailHSM uses the ADSM NetWare API to communicate with ADSM. The NetWare server operates as an ADSM client to the ADSM server. You can run AvailHSM on NetWare 3.x and 4.x servers. You can use any ADSM server that supports TCP/IP, IPX/SPX, or APPC.

AvailHSM

AvailHSM automatically migrates older, infrequently accessed data from the NetWare server volumes to lower-cost media within the storage hierarchy for NetWare servers. The data moves from primary to secondary storage volumes. You see the migrated data as if it resides on a primary storage volume. The migration process gives you the impression of having unlimited disk space. If you read or update migrated data, the data is automatically recalled from AvailHSM to the NetWare volume from where it was originally migrated.

A typical AvailHSM configuration consists of an AvailHSM storage server and a number of managed NetWare servers. The storage server controls the migration and recall of data from and to the managed servers and the movement of data in the storage hierarchy.

The first level of media in the secondary storage hierarchy is always defined on disk. Up to two additional storage levels can be defined beneath this disk level. The additional storage levels can be either optical or tape devices locally

attached to the AvailHSM server. To use ADSM with AvailHSM, you define a logical device for one of the additional storage levels.

AvailHSM and ADSM

The integration of AvailHSM and ADSM provides a complete solution for storage management in a NetWare environment. AvailHSM provides storage management services, and ADSM provides the backup and recovery services as well as being the data repository for AvailHSM.

A virtual device driver allows ADSM to be defined to AvailHSM as a virtual tape library. This addition to AvailHSM, called the AvailHSM/ADSM Virtual Tape Library Storage Capacity Kit (ADSM Capacity Kit), enables an ADSM server to become a level in the AvailHSM secondary storage hierarchy. The ADSM Capacity Kit uses the ADSM NetWare client API to communicate with an ADSM server. It is the link between AvailHSM and the ADSM server.

AvailHSM sees the ADSM server as a tape library device, and the ADSM server sees AvailHSM as an additional application using its storage pools.

Order Information

The ADSM API and the ADSM Capacity Kit are provided by Wang Software Storage Management Group, Inc. U.S customers can order them from IBM. Other customers can order them from a Wang reseller.

More Information

For more information about AvailHSM, see:
http://www.wang.com/sbu/smg/wanghsm.htm and *Hierarchical Storage Management for NetWare: ADSM and AvailHSM Implementation* (SG24-4713).
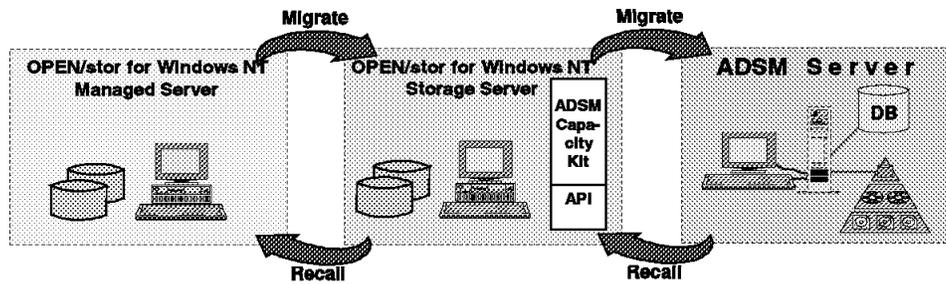
OPEN/stor for Windows NT

The ADSM Windows NT client supports the backup-archive client only. ADSM does not provide an HSM solution for the Windows NT environment. The ADSM space management solution for Windows NT servers and workstations is OPEN/stor Hierarchical Storage Management for NT (OPEN/stor). OPEN/stor is available as a separate product from Wang Software Storage Management Group, Inc. It complements the backup and archive functions provided by the ADSM Windows NT client. OPEN/stor uses the ADSM Windows NT API to communicate with ADSM. The Windows NT server or workstation operates as an ADSM client to the ADSM server.

You can run OPEN/stor on Windows NT 3.51 and 4.0. You can use any ADSM server that supports TCP/IP, IPX/SPX, or APPC.

OPEN/stor

OPEN/stor automatically migrates older, infrequently accessed data from the Windows NT file systems to lower-cost media within the storage hierarchy for Windows NT. The data moves from primary to secondary storage volumes. You see the migrated data as if it resides on a primary storage volume. The migration process gives you the impression of having unlimited disk space. If you read or update migrated data, the data is automatically recalled from OPEN/stor to the Windows NT file system from where it was originally migrated.

A typical OPEN/stor configuration consists of an OPEN/stor storage server or workstation with attached magnetic, optical, and tape storage devices. The storage server or workstation controls the migration and recall of data from and to the managed servers and the movement of data in the storage hierarchy.

The first level of media in the secondary storage hierarchy is always defined on disk. Up to two additional storage levels can be defined beneath this disk level. The additional storage levels can be either optical or tape devices locally attached to Windows NT. To use ADSM with OPEN/stor, you define a logical device for one of these additional storage levels.

OPEN/stor and ADSM

The integration of OPEN/stor and ADSM provides a complete solution for storage management in a Windows NT environment. OPEN/stor provides storage management services, and ADSM provides the backup and recovery services as well as being the data repository for OPEN/stor.

A virtual device driver allows ADSM to be defined to OPEN/stor as a virtual tape library. This addition to OPEN/stor, called the OPEN/stor/ADSM Virtual Tape Library Storage Capacity Kit (ADSM Capacity Kit), enables an ADSM server to become a level in the OPEN/stor secondary storage hierarchy. The ADSM Capacity Kit uses the ADSM Windows NT client API to communicate with an ADSM server. It is the link between OPEN/stor and the ADSM server.

OPEN/stor sees the ADSM server as a tape library device, and the ADSM server sees OPEN/stor as an additional application using its storage pools.

Order Information

The ADSM API and the ADSM Capacity Kit are provided by Wang Software Storage Management Group, Inc. U.S customers can order them from IBM. Other customers can order them from a Wang reseller.

More Information

For more information about OPEN/stor, see http://www.wang.com/sbu/smg/wanghsm.htm.
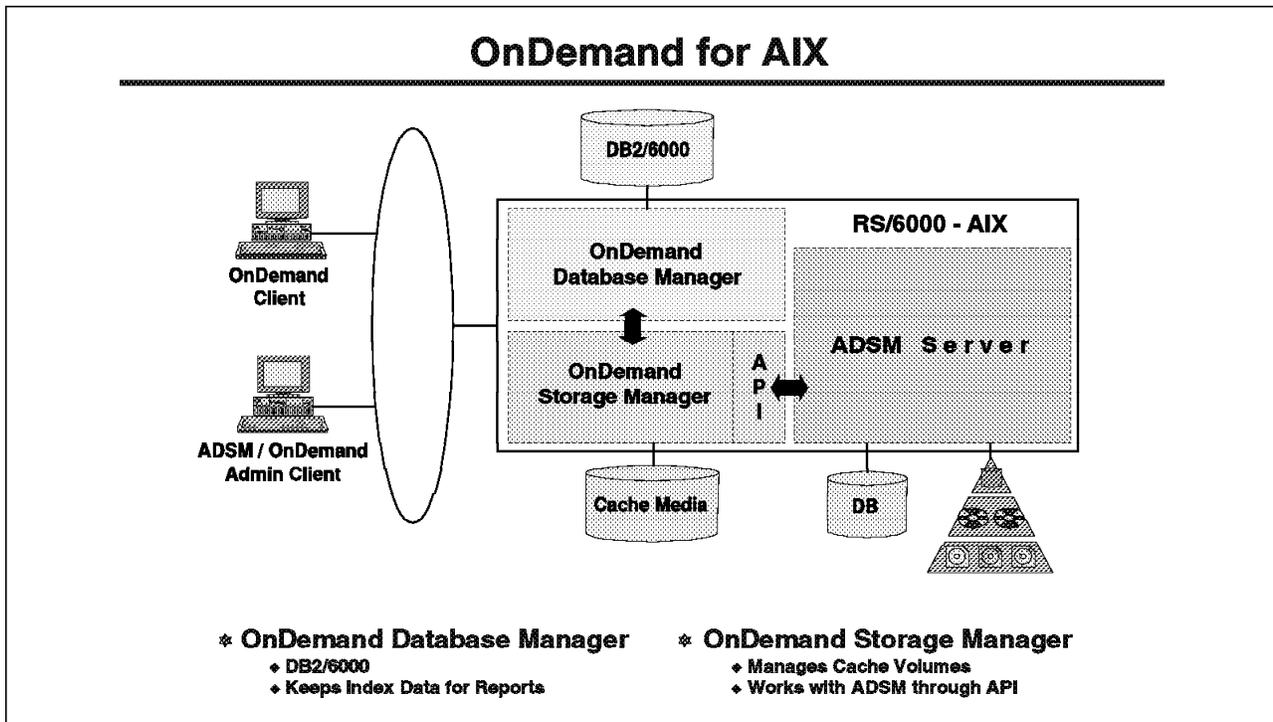
OnDemand for AIX

OnDemand for AIX

IBM's OnDemand for AIX is a client/server archival and retrieval system for documents that offers you an alternative to microfiche and hardcopy storage.

OnDemand for AIX provides a set of programs to capture, index, store, and retrieve large volumes of computer-generated documents. Once documents are indexed and stored, users running the OnDemand client GUI under Windows or OS/2 can search, retrieve, view, print, and fax documents.

An OnDemand system consists of:

- Client programs
- Server programs
- A network running TCP/IP
- Database management programs
- Storage management programs

The client programs are your gateway to the report files archived in OnDemand. You can construct queries and search for report files; retrieve items from OnDemand; view, print, and fax copies of report files; and attach electronic notes to report file pages.

The server programs manage control information and index data, store and retrieve report files, and process query requests from the client programs. Report files can reside on disk, optical, and tape storage volumes.

Client and server programs communicate over a network running TCP/IP.

OnDemand Database Manager

OnDemand uses DB2 for AIX as the database manager to maintain index data for the report files that you store in OnDemand. The database manager maintains the OnDemand database with tables that describe the applications, application groups, storage sets, folders, groups, users, and printers that you define.

OnDemand Storage Manager

OnDemand provides a high-speed cache outside ADSM. The storage manager maintains files on cache storage volumes and works with ADSM to manage files stored on archive media, such as optical or tape storage libraries. ADSM stores and retrieves files from archive media. In general, OnDemand should use ADSM only for managing storage pools of optical and tape devices. When OnDemand loads a report file, it can simultaneously write the file to cache volumes and ADSM-managed optical or tape storage. Once you have decided which types of media your OnDemand system supports, you configure the storage device and define the storage device volumes to the ADSM server.

OnDemand communicates with ADSM for AIX through the ADSM AIX client API.

Order Information

You can order OnDemand for AIX (product number 5622-662) from IBM.

More Information

For more information about OnDemand for AIX, see *OnDemand for AIX: Introduction and Planning Guide* (G544-5281).

Backup Recovery and Media Services/400

* **Backup and Recover**
* **Archive and Retrieve**
* **Media Management**
* **Guided Recovery**
* **BRMS/400 and ADSM**
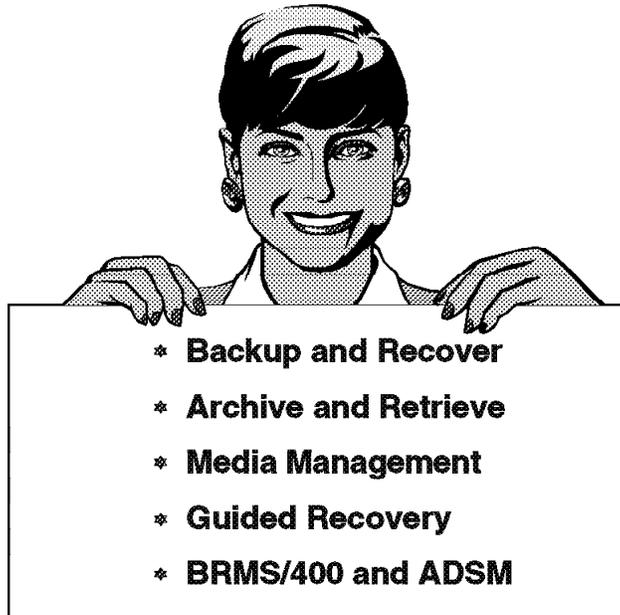
Backup Recovery and Media Services/400

ADSM provides storage management for distributed data, and IBM's Backup Recovery and Media Services/400 (BRMS/400) provides storage management for your AS/400 system.

Backup and Recover

You can back up systems, libraries, files, documents, and spool files for your AS/400 system. Backup control groups enable you to use a single command to back up and recover different types of files with common characteristics as an entity. You can use a job scheduler to automate your backups.

Archive and Retrieve

The archive function enables you to copy files to offline media. When the copy has been written to the media, the file is deleted from the disk.

Media Management

BRMS/400 also provides you with facilities to manage your media, such as enroll new, initialize, move, expire, and duplicate.

When you move a tape volume from one location to another, BRMS/400 keeps track of the location.

If you have multiple AS/400 systems, you can have a single tape pool for all of them.

Guided Recovery

BRMS/400 can guide you through a disaster recovery with a list that includes all necessary recovery steps. In case of a total system failure, BRMS/400 uses this list to make the recovery.

BRMS/400 and ADSM

BRMS/400 and ADSM provide similar functions, but remember that BRMS/400 manages files on the AS/400 system, whereas ADSM manages files for distributed data in desktop and UNIX machines connected to your ADSM server.

Order Information

You can order BRMS/400 from IBM. The product number depends on the OS/400 version and release you have installed.

LAN Services for OS/390 and ADSM

IBM's LAN Services for OS/390 includes elements of LAN Resource and Extension Services (LANRES) and LAN Server, both of which enable your LAN servers to take advantage of the availability, performance, security, and administration of MVS.

LANRES and LAN Server

LANRES provides disk serving, print serving, data distribution, and central administration for NetWare LAN clients connected to MVS.

LAN Server provides file serving to OS/2 and Network File System (NFS) servers connected to S/390 servers. NFS servers can be AIX or UNIX. LAN Server also provides the capability for OS/2 and NFS clients to share a common data repository with full update capability.

LAN Services for OS/390 provides rapid communication between MVS and the LAN and is transparent to the LAN clients. The LAN servers can communicate with MVS through TCP/IP or ESCON- or parallel-attached channels. LAN servers can also use APPC as the communication protocol. The MVS system becomes a high-speed file server and allows data to be stored in a VSAM data set, which can be shared among all LAN servers. Thus disk space can be freed up on the workstation-based LAN servers, and the large capacity of MVS can relieve the constraints of those LAN servers. Users access files on the host system as if the files were on their LAN server or local disk.

LAN Services and ADSM

ADSM clients back up their files to the ADSM storage pools, regardless of whether the files reside on local disk or are a part of the VSAM data set. The VSAM data set itself is space managed by DFSMS, just like every other MVS data set. Without ADSM you can back up only the VSAM data set; with ADSM you can back up individual files within the data set.

Support for Other Platforms

LANRES is also supported on the VM and VSE platforms. The VM and VSE product corresponding to the LAN Server is called the LAN File Server (LFS) for VM and VSE.
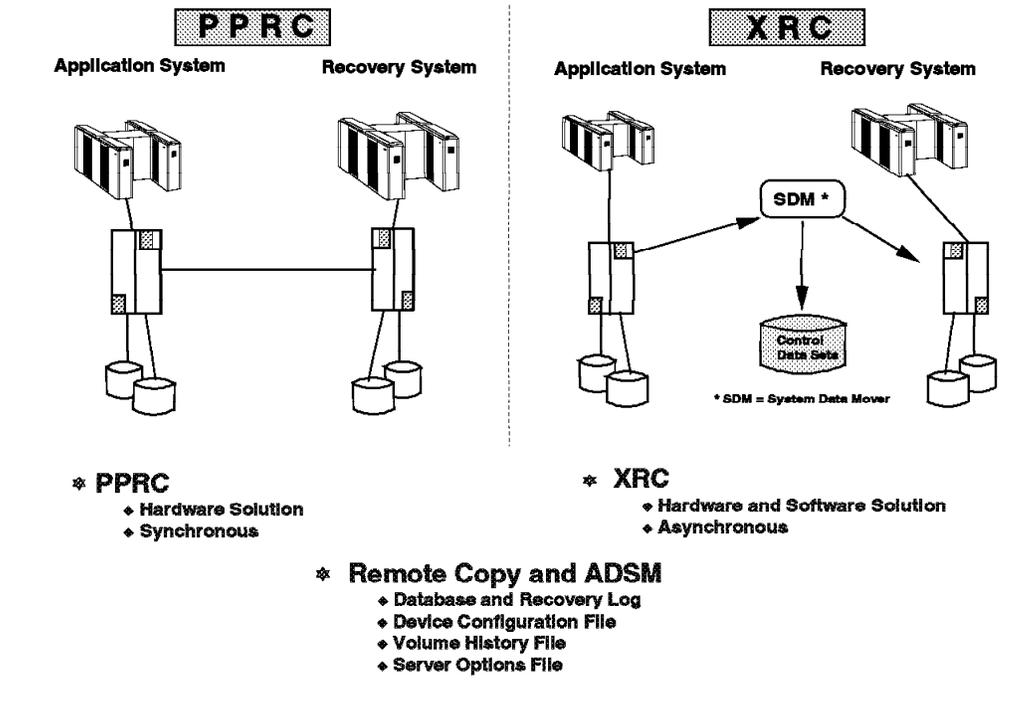
Remote Copy and ADSM

Remote copy is a storage-based disaster recovery and workload migration solution that enables you to copy disk data in real time to a remote location. It is a function of the IBM 3990 Model 6 Storage Control, along with the appropriate levels of MVS/DFP and DFSMS/MVS.

Remote copy offers two options:

- Peer-to-peer remote copy (PPRC)

- Extended remote copy (XRC)

Both options address the problem of unrecoverable data that occurs between the last, safe backup of an application system to a recovery system and the time when the application system fails.

Disasters occur in many forms. Some occur suddenly and stop all processing at a single point in time; more frequently, others interrupt operations in stages, occurring over several seconds or minutes. The remote copy function addresses the problem caused by intermittent and gradual system failures.

Remote copy is based on two systems, an application system and a recovery system, either in the same building or at remote locations. Each system has specific data that you have identified as remote-copy-managed. Data is written to the primary disk and shadowed to the secondary disk transparently to application users.

PPRC

PPRC is a hardware solution designed for those sites that must have data at the recovery system fully current with the application system and can accept some performance impact on the write operations of applications.

With PPRC, no disk data is lost between the last update at the application system and the recovery at the recovery system.

A PPRC data copy to the recovery system is synchronous with the primary volume's write operations. PPRC does not consider the application system disk write complete until the recovery system write is complete. Each application system write to the recovery system causes an increase in the response time of the application system.

The maximum recommended distance between the application system and the recovery system is 43 km (27 mi).

XRC

XRC is a combined hardware and software solution designed for those sites that have to maintain the highest levels of performance on their application system and can accept a gap of a few seconds between writes on the application system and writes on the recovery system. There may be little or no data loss if a system failure occurs because the currency of data at the recovery system is only seconds behind the primary write operation.

With XRC, the only data lost is data in transit between the time the application system fails and the recovery system recovers.

XRC allows the primary application's write operation to signal completion before the data is shadowed to the secondary volume.

The maximum distance between the application system and the recovery system is unlimited if channel extenders are used.

For more information about channel extenders, see "CNT Channelink and ADSM" on page 194.

Remote Copy and ADSM

If you decide to manage your ADSM for MVS disk data through PPRC or XRC, the data is automatically written to the primary as well as the secondary system in real time.

We recommend that you use remote copy to manage the ADSM database, recovery log, device configuration file, volume history file, and server options file.

More Information

For more information about remote copy, see the *Planning for IBM Remote Copy* (SG24-2595), *P/DAS and Enhancement to the IBM 3990-6 and RAMAC Array Family* (SG24-4724), and *Remote Copy Administrator's Guide and Reference* (SC35-0169).

CNT Channelink and ADSM

The Channelink family from Computer Network Technology (CNT) provides high-performance, long-distance, networked channel extension for a complete range of mainframes, distributed systems, and storage devices.

With Channelink and ADSM you get the following facilities:
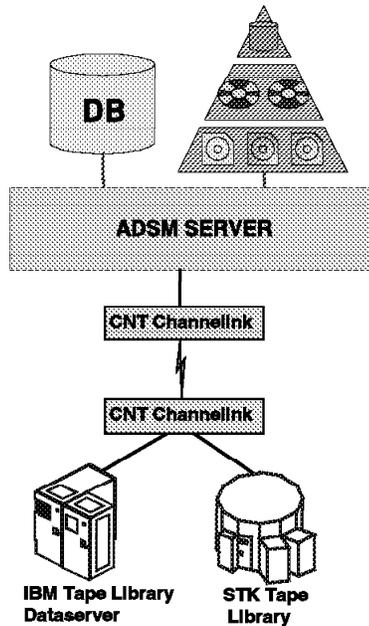
- Remote vaulting for disaster recovery
- Elimination of channel distance restrictions
- Expanded interoperability between ADSM servers and storage devices
- Software and hardware support

Remote Vaulting for Disaster Recovery

ADSM and Channelink enable you to direct the primary storage pools to a remote location. You can also direct your database copy and copy storage pools to a remote location. It is also advisable to keep a copy of the volume history, device configuration, and server options files at the remote location.

The data is automatically moved by the Channelink, so you do not have to manage the movement. Channelink uses telecommunication lines to transfer the data.

Elimination of Channel Distance Restrictions

Parallel channels are limited to distances of about 120m (400 ft). With ESCON channels you can extend the limit to 43 km (27 mi) for DASD. With SCSI interfaces the limit varies between 6 and 25 m (20 - 80 ft).

With the Channelink channel extension for MVS and the SCSI Gateway for UNIX, you can access storage devices at a distance of up to 1600 km (1000 mi).

Expanded Interoperability between ADSM Servers and Storage Devices

In both MVS and UNIX environments, Channelink products can be used to share storage resources between ADSM servers. Therefore you can use the devices efficiently regardless of their location.

Software and Hardware Support

Channelink is based on a family of network processors, software products, and services provided by CNT. Each network processor supports many combinations of devices and network interfaces.

Channelink supports ADSM servers for MVS, AIX, HP-UX, and Sun Solaris. Tape support is for all types of 3480s, 3490s, and 3590s, including the 3494 and 3495 Tape Library Dataservers and STK tape libraries. Channel connectivity support includes ESCON, bus-and-tag, and SCSI-2 Fast and Wide.

More Information

For more information about CNT's Channelink, see http://www.cnt.com.

ADSM for AIX StorageTek Tape Library Support

ADSM for AIX StorageTek Tape Library Support enables you to extend the use of automated operations to the LAN through ADSM and the StorageTek tape libraries. Tape libraries can be effectively shared among RS/6000 or SP2 servers, or with an MVS host. The provided tape drivers support StorageTek tape drives in native mode and do not emulate IBM 3490 tape drives. Support of StorageTek tape drives in native mode allows the use of mixed tape media in the same library. The reporting system in the tape libraries keeps track of the different drive types.

RPQ

The RPQ includes a set of utilities and features that interact with StorageTek's ACSLS server or Library Station software. The utilities query the status of volumes, drives, and scratch pools; allow for mounting and dismounting of tape volumes; and enable the definition of scratch pools and the assignment of tape volumes to scratch pools. You can check a large number of tape volumes into the library at one time.

Device Support

The following tape libraries are supported:

- 4410 LSM
- 9310 Powderhorn
- 9360 Wolfcreek
- 9710 Timberwolf

LSM, Powderhorn, and Wolfcreek tape libraries can contain the following control units:

- 4781 18-track through a SCSI connection

- 4791 36-track through a SCSI connection

- 9490 Timberline through a SCSI-2 Fast and Wide connection

- SD-3 Redwood through a SCSI-2 Fast and Wide connection

The Timberwolf tape library can contain 4890 Twinpeaks and/or Quantum DLT 400 tape drives.

Order Information

You can order the RPQ from IBM; the product number is 5799-XQW.

Complementary Products for Database Backup

This graphic lists solutions that work with ADSM to provide database backup.

Backup Restore Interface for Oracle for SAP R/3

IBM Backup Restore Interface for Oracle for SAP R/3 (BACKINT/ADSM) is an interface program that connects the SAP R/3 database administration utilities with ADSM. This solution is designed to shorten the backup and restore time to minimize the impact of system management tasks on the availability of the SAP applications.

For more information, see "ADSM and Oracle" on page 219 and http://www.ibm.de/fo/ide/solutions/r3_adsm/r3_adsm.html.

Backup and Restore Interface for Informix (ON-Bar)

This X/Open BSA-compliant utility allows customers to back up Informix databases directly to ADSM for the AIX and Sun Solaris platforms. Support is provided through the ADSM X/Open API client.

For more information, see "ADSM and INFORMIX-OnLine" on page 231.

ADSMConnect Agent for Oracle Backup

The ADSMConnect Agent for Oracle Backup is an interface between the Oracle Enterprise Backup Utility (EBU) and the ADSM API for the AIX client platform. The ADSMConnect agent enables ADSM to back up Oracle7 databases online by using the Oracle Enterprise Backup Utility.

For more information, see "ADSM and Oracle" on page 219.

SQL-BackTrack for Oracle or Sybase

DataTool's SQL-BackTrack ADSM module for Oracle or SQL BackTrack ADSM module for Sybase integrates the database-aware backup and recovery features of SQL-BackTrack with the services of ADSM. The ADSM module sends SQL-BackTrack backups directly to an ADSM server, so you can build an enterprisewide backup and recovery solution for databases and file systems.

For more information, see "ADSM and Oracle" on page 219, "ADSM and Sybase" on page 226, and http://www.datatools.com.

Lotus Notes Backup Agent

The ADSM Lotus Notes Backup Agent is an integrated Lotus Notes application running on OS/2. Support for Lotus Notes databases on Windows NT and AIX is a high priority requirement. It performs incremental and full backups of OS/2 Notes databases and the Lotus Notes and ADSM API to back up and restore OS/2 notes documents and databases to and from an ADSM server.

For more information, see "ADSM and Lotus Notes" on page 240.

Chapter 12. ADSM and Databases
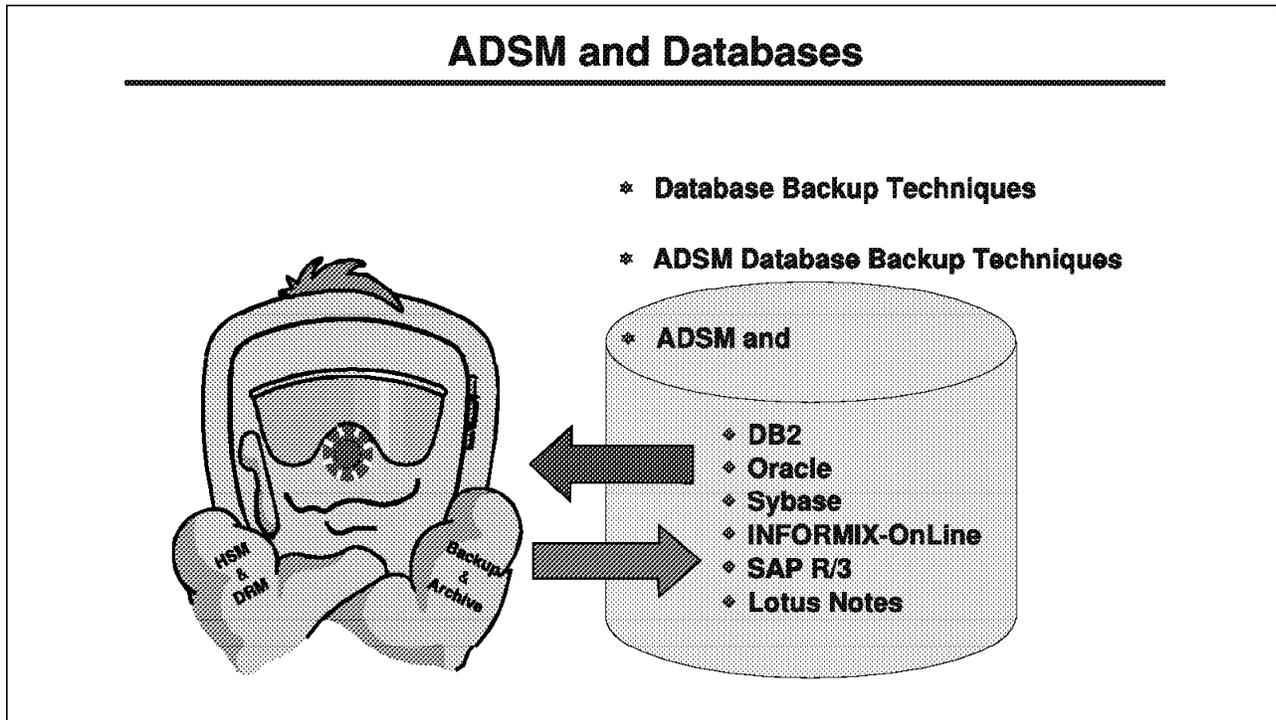
ADSM and Databases

ADSM can be used to back up workstation databases.

This chapter describes:

- Database Backup Techniques

- ADSM Database Backup Techniques

- ADSM and DB2/6000, DB2/2, Oracle, SAP R/3, Sybase, INFORMIX-OnLine, Lotus Notes

Most of this material was taken directly from the redbook: *Using ADSM to Back Up Databases* (SG24-4335) For more details on backing up these and other databases with ADSM, refer to the redbook.

Database Backup Techniques
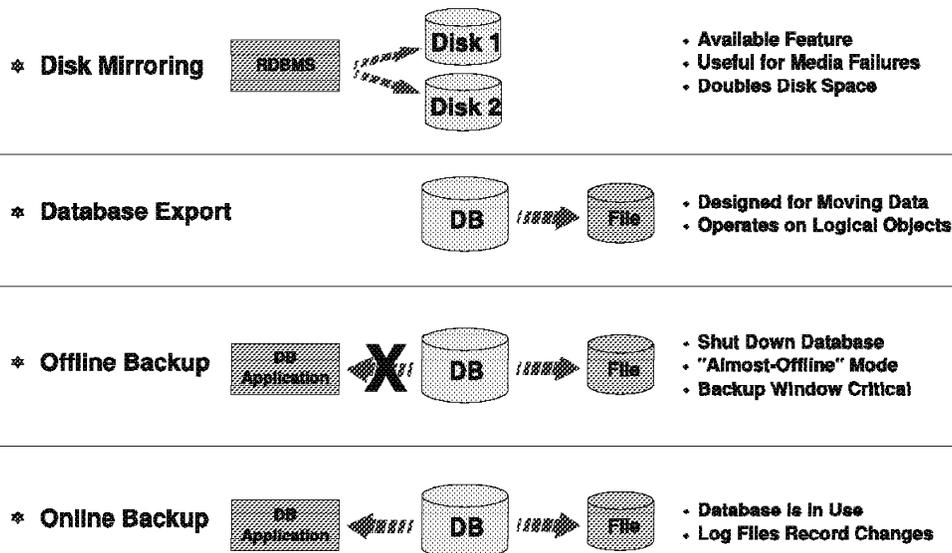
Database Backup Techniques

You can use a number of techniques to back up data managed by relational database management systems (RDBMSs). These techniques are, at least at a conceptual level, common to most RDBMSs.

Disk Mirroring

Disk mirroring is a useful technique for maximizing the availability of your database. Mirroring is the process of writing the same data to multiple storage devices at the same time. If a media failure occurs, operations are automatically switched to a mirrored copy. Mirroring allows your users to continue working even though a media failure has occurred. Mirroring can be implemented in either software or hardware; for example, AIX, most RDBMSs, and the IBM 3990 Model 3 storage controller provide mirroring facilities.

However, mirroring does not remove the need to back up databases. For example, disk mirroring will not allow you to restore a table that has been lost or damaged as a result of user error. Also, although disk mirroring dramatically reduces the impact of media failures, there is still a risk of damage to both sides of the mirror. If a database is held on one set of physical volumes, and a mirror image of the same database is maintained on a separate set of physical volumes, it is possible for both sets of physical volumes to be damaged or destroyed. This could happen as a result of a disaster or it could just be bad luck. In such instances, it will be necessary to recover the database from backup copies.

Database Export

All RDBMSs provide export and import utilities. These utilities operate on logical objects as opposed to physical objects. For example, you can use an export command to copy an individual table to a file system file. At some later time, you might want to restore the table, in which case you would use the import command. Although export and import can be used for backup and restore operations, they are really designed for moving data, for example, for workload balancing or migration.

Note: You should assume that import will work only with files that have been created by the same RDBMS's export utility.

Most other non-RDBMS utilities (such as operating system utilities and products like ADSM) operate on the physical data files that RDBMSs use to store their databases. Therefore, other utilities cannot usually be used to back up and restore a single table because:

- A single physical data file may contain data belonging to several tables.

- The data in a single table may be spread across multiple data files.

Thus, the only way to gain access to the set of data in a single table is through the RDBMS itself.

Export utilities are usually slower than most other utilities and should be used only when you need access to database objects or raw devices.

Offline Backup

Offline backup involves shutting down the database before you start the backup and restarting the database after backup is complete.

Offline backups are relatively simple to administer. However, they suffer from the obvious but significant disadvantage that neither users nor batch processes can access the database while the backup is taking place. You must schedule sufficient time to perform the backup to ensure that the periods when the database will be unavailable are acceptable to your users.

Some RDBMSs provide a "single-user mode" or "quiesced mode." You can think of this as an "almost offline" mode. A database administrator can still use the database, but general users cannot.

Online Backup

Some (but not all) RDBMSs allow backups to be performed while the database is started and in use.

Clearly, if a database is being backed up while users are updating it, it is likely that the data backed up will be inconsistent ("fuzzy"). The RDBMSs that support online backup use log files during the recovery process to recover the database to a fully consistent state. This approach requires that you retain the RDBMS log files and indicate to the RDBMS when you are about to start the backup and when you have completed the backup.

Some RDBMSs allow you to quiesce activity on portions of the database (for example, a particular table space) so that a set of complete tables is temporarily "frozen" in a consistent state. You then can back up the set of tables that has

been "frozen." Once the backup has completed, you can reactivate the table space.
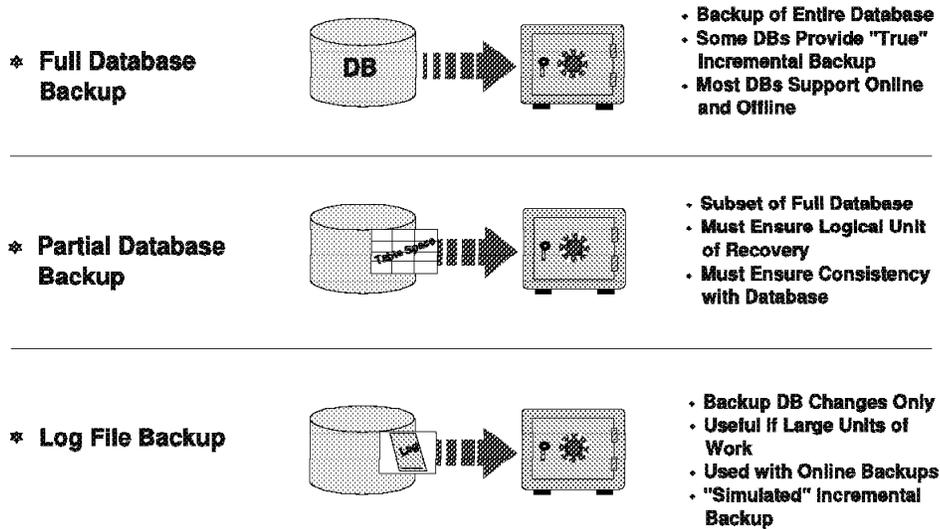
Database Backup Techniques...

Full Database Backup

Full database backups involve making copies of all of the data files used to hold user data. In some database products, full database backups also include copies of the data files that hold tables used by the RDBMS itself, RDBMS log files, and any control files and parameter files that the RDBMS uses. Many RDBMSs allow you to perform full database backup when the database is either online or offline.

Some database products provide an incremental backup that only backs up changed database pages or blocks. This type of incremental backup is called a "true" incremental backup, as opposed to a "simulated" incremental backup (see log file backup).

Partial Database Backup

Partial database backups involve backing up a subset of the full database (such as the data files that make up a table space). Make sure that the subset you back up (as part of a partial backup) represents a complete logical unit of recovery from the point of view of the application. You may also need to back up data files that the RDBMS does not manage.

You must also ensure that the unit of recovery is consistent from the point of view of the RDBMS. If you have added a new data file to a table space, you must ensure that any control file that the RDBMS uses to define the relationship between data files and table spaces is also backed up.

Many RDBMSs allow partial database backups when the database is online or offline.

Log File Backup

For some applications, the units of recovery are too large to be backed up on a daily basis. Sometimes the constraining factor is the elapsed time that is available (the backup window). Sometimes the load that the backup would place on the network would have an unacceptably bad impact on other processes and users.

In such situations it may be possible to capture only the changes to the database by backing up the RDBMS′ log files. This type of backup is sometimes referred to as a "simulated" incremental backup. Recovery in such situations is achieved by:

1. Restoring the database from a full database backup (in some circumstances restoring from a partial backup may be sufficient)

2. Restoring the log files

3. Applying the log files to the restored database.
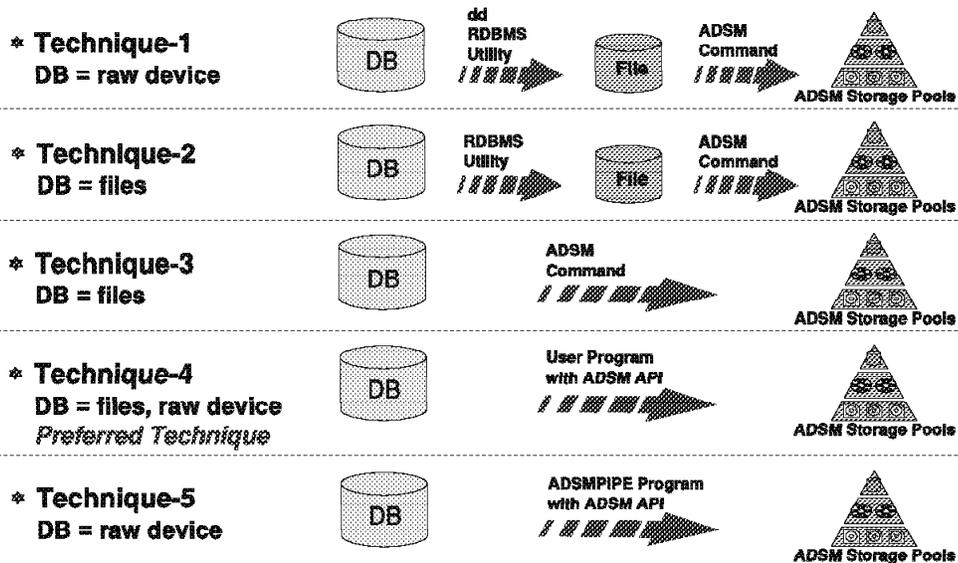
ADSM Database Backup Techniques

You can use any of several techniques to back up databases with ADSM. As shown on the graphic, the techniques can involve using operating system utilities, RDBMS utilities, ADSM, or a combination of a utility with ADSM. They are applicable to RDBMSs on UNIX as well as on OS/2, with one exception. Technique-5 applies only to implementations on raw devices, which are unknown in OS/2.

Technique-4 is the preferred method and should be used whenever possible. Use the other techniques only when a program that uses the ADSM API is not available.

Technique-1

Technique-1 applies to databases installed on raw devices. If the database or application has not been enhanced to use the ADSM API, then, in order to use ADSM to back up a database installed on raw devices, you must convert the raw devices to one or more files. You can use the *dd* command for this conversion, but the command will back up the entire contents of the raw device, and your database must be offline. You can use RDBMS utilities to convert the raw devices, which will allow backup at the RDBMS object level and, depending on the database product, might allow online backup.

Whether you use the dd command or RDBMS utilities for the conversion, you must be sure that you have sufficient disk space available to hold the files that they create. If the RDBMS utilities provide backup at an object level, such as a table space, you might be able to reduce the amount of additional disk space you would need for the files. Remember also that, when you restore the data that ADSM has backed up, you will need sufficient disk space to hold the files.

Technique-2

Technique-2 applies to databases installed on files. You can choose to use an RDBMS utility to back up the data to files and then use ADSM to back up those files. Use this technique if you want to back up something smaller than the entire database, such as a table space, or if you want to take advantage of an RDBMS utility that allows online backups. The same consideration mentioned for Technique-1 still applies: ensure that you have enough disk space.

Technique-3

Technique-3 also applies to databases installed on files. With this technique, you use ADSM to directly back up the files that make up the database and logs. The advantage of this technique is that intermediate files are not created, but the database must be offline to ensure a consistent copy. The exception to this is Oracle. You can use ADSM to directly backup Oracle files while the database is online by using the alter tablespace commands (see "Alter Tablespace" on page 220).

Technique-4

Technique-4 applies to databases installed on either raw devices or files. If the application or database product has been enhanced to use the ADSM API, the application handles the underlying physical structure, so whether raw devices or files are used is irrelevant. Also, the application determines and controls the type of backup (for example, online, offline, at the table level).

Technique-5

The ADSMPIPE utility is provided as an "as-is" utility with the redbook *Using ADSM to Back Up Databases (SG24-4335)* and on the ITSO ftp server: ftp.almaden.ibm.com (in the SG244335 directory). ADSMPIPE aids in the backup of databases installed on raw devices. It reads raw devices and sends the data to ADSM using the ADSM API. It is an offline, full backup alternative. There is one exception where you can produce an online backup, Oracle!

ADSMPIPE is similar to using the dd command (technique-1) except that an intermediate file is not created.

Use ADSMPIPE *only* if a better API alternative is not available.
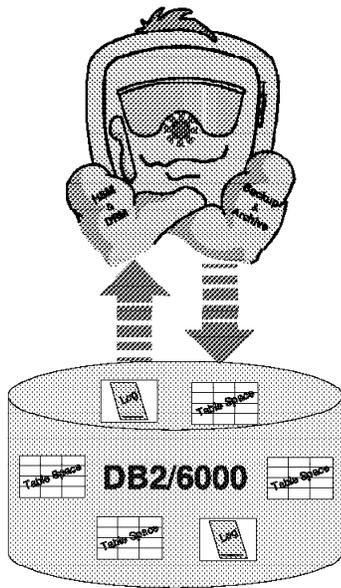
ADSM and DB2/6000

DB2 support of the ADSM API is provided for following operating system platforms:

- **AIX**
- **OS/2**
- **Sun Solaris**
- **HP-UX**
- **Windows NT**
- **Sinix**

In this section we describe the DB2 support for AIX and OS/2.

DB2/6000 Structure

The key components of a DB2/6000 database that you must consider for backup and recovery are the:

Database
The database contains both control information and user data. The system catalog tables contain the control information. They describe the logical and physical structure of the data.

Configuration Files

The configuration files contain parameter values that define the resources allocated to DB2/6000 and individual databases.

Directories

The directories are used to access local and remote databases.

Recovery Logs

All databases have recovery logs associated with them that keep records of database changes. In DB2/6000 the active logs prevent a failure from leaving the database in an inconsistent state. They contain information for transactions whose changes have not yet been written to the database files. In case of a failure, the changes already made but not committed are rolled back, and all committed transactions, which may not have been physically written to disk, are redone. These actions ensure the integrity of the database.

Roll-forward recovery uses logs to allow a database to be rebuilt to a specified point in time. In addition to using the information in the active logs to rebuild a database, the roll-forward function uses archived logs to reapply previous changes.

DB2/6000 Version 1 Backup Utilities

Backup/Recover

DB2/6000 Version 1 provides a utility called *backup* for online and offline backup of an entire database.

The online backup is done while the database is online. Thus users or applications can connect to the database and make changes while the backup is running. The database image file is not consistent. Therefore it is not possible to use the database image file alone to perform a recovery operation because active units of work are present in both log files and database tables.

For offline backup, no other processes can be connected to the database at the same time. Each application must be logged off. When the backup operation starts, SQL operations are not permitted.

The DB2/6000 Version 1 recovery utility rebuilds a database from a previously saved database backup image. It is an offline utility. There are two recovery methods:

- **Restore recovery**, which rebuilds a database to its state when the backup copy was made. You can also use this method to duplicate a database.

- **Roll-forward recovery**, which builds a database from a restored copy, making changes to it since the time the backup was made. These changes are found in the log files.

The DB2/6000 backup utilities are integrated with ADSM services by using the ADSM API. Thus an intermediate file is not created during the backup operation before the database image is stored on the ADSM server. Both online and offline backups can be performed with ADSM.

Export/Import

DB2/6000 also provides an export/import utility to move data. This utility can be used as a supplement to your backup strategy, but it is not a replacement for the backup utility. There is a risk of introducing inconsistencies into your database because no synchronization of data with logs is performed. However, the DB2/6000 Version 1 backup utility only backs up entire databases, so use of the export/import utility is the only way to capture individual tables.

DB2/6000 Version 2 Enhancements

Raw Device Support

DB2/6000 Version 2 can (unlike Version 1) store data directly on a raw device and thus gives database administrators the choice of implementing databases on raw devices or file systems.

Table Space Concept

With DB2/6000 Version 2, databases can be partitioned into table spaces.

Backup/Recovery Enhancements

DB2/6000 Version 2 provides significant enhancements to the backup/restore utilities. Backup and recovery can be performed at the table space level. If a table space contains a single table, the backup or recovery is equivalent to a table level backup or recovery. Table spaces can be online or offline during the backup process. During recovery, all table spaces, other than the one being recovered, can remain online.

Another enhancement to the backup/recovery process is *parallel backup and recovery*. It is possible to perform the backup or recovery of a database or table space in parallel to or from multiple devices. This enhancement can dramatically reduce the elapsed time requirements.

New Load Utility

The load utility significantly increases the speed of doing data loads and ensures recoverability of the data being loaded. It is intended for either bulk loading of new tables or appending large amounts of data to existing tables. The load utility is restartable and recoverable. If a failure occurs while loading data, you can continue the load without starting from the beginning. To ensure recoverability, a backup copy of the loaded table is created during the load processing.

DB2/6000 Backup Alternatives

Backup Type	Database Implementation	Database Mode	Intermediate File Created	Backup Granularity	Platforms Supported
DB2/6000 V1 Backup	File System	Online, Offline	No	Database	AIX
DB2/6000 V2 Backup	File system, Raw Device	Online, Offline	No	Database, Table Space	AIX HP Sun Sinix
DB2/6000 V1 and V2 Export	File System, Raw Device	Online, Offline	Yes	Table	AIX HP Sun Sinix

DB2/6000 Backup Alternatives

This graphic gives an overview of the DB2/6000 backup alternatives with ADSM and summarizes some of their main characteristics, such as:

- Permitted ways of implementation (on file systems or raw devices)

- Permitted database modes (online or offline)

- Creation of an intermediate file (yes or no)

- Backup granularity (the entities that can be backed up)

The DB2/6000 support of the ADSM API is provided for AIX, Sun, HP, Sinix, and NT platforms.

DB2/6000 Parallel Edition

A parallel database must "push the envelope" on performance and capacity; therefore it must fully leverage the capabilities of the underlying hardware, operating system, and computing environment.

Parallel Database Architecture DB2/6000 Parallel Edition (DB2PE)

DB2/6000 Parallel Edition (DB2PE) is a key member of the DB2 family of database engines. Parallel database technology has been under development at the Almaden Research Center and elsewhere in IBM for several years.

Shared-Nothing

DB2PE exploits a shared-nothing architecture, where each node has its own processor, memory, and disk, and effectively "nothing" is shared. A shared-nothing architecture has the best scaleability because most of the processing is done independently on each node, with minimal interaction among nodes.

Database Stored across a Network of Processors

Each node in a DB2PE system supports a subset of the overall database. For a given user request, data access, buffer management, locking, logging, and other processing is done, where possible, independently on each node in parallel. The database appears to the application user as a single database on a single system.

Linearly Scaleable

DB2PE exploits the IBM POWERParallel SP2. However, it is in no way tied to or dependent on that hardware platform. DB2PE can run as multiple instances on a RS/6000 uniprocessor as a means of doing parallel I/O. On an SMP model of a RS/6000, multiple DB2PE instances can be used to leverage all processors in the single machine. DB2PE can also run on separate RS/6000s connected together on a LAN. Different RS/6000 models can be used, including a 2- or 4-way HACMP configuration.

Backup Tools Can Back Up from Each Node

DB2PE is based on DB2/6000 Version 1; it has the same level of function as DB2/6000, including its backup and restore functions. DB2PE Version 1 provides full online and offline backup and uses the ADSM API so that backups can be sent directly to ADSM. To back up a DB2PE database with ADSM, you must back up each individual DB2PE node.

A detailed redbook on using ADSM to back up DB2PE is available: *Backup, Recovery, and Availability with DB2 Parallel Edition on RS/6000 SP* (SG24-4695).

Other Products Using Parallel Database Architecture

Two other database products use a parallel database architecture, INFORMIX-OnLine Extended Parallel Server (XPS) and Sybase.

For more information about INFORMIX-OnLine XPS refer to "INFORMIX-OnLine Extended Parallel Server (XPS)" on page 235.
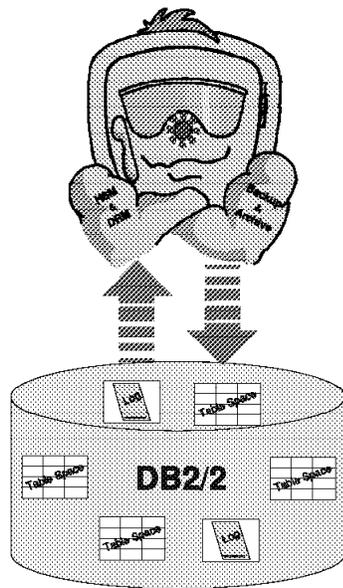
ADSM and DB2/2

When we refer to DB2/2 in this section, we are referring to DB/2 on OS/2 and Windows NT.

DB2/2 Structure

The key components of a DB2/2 database that you must consider for backup and recovery are the:

Database

The database contains both control information and user data. The control information is contained in system catalog tables. These tables describe the logical and physical structure of the data.

Configuration Files

The configuration files contain parameter values that define the resources allocated to DB2/2 and individual databases.

Directories

The directories are used to access local and remote databases.

Recovery Logs

All databases have recovery logs associated with them that keep records of database changes. In DB2/2, the active logs prevent a failure from leaving the database in an inconsistent state. These logs contain information for transactions whose changes have not yet been written to the database files. In case of a failure, the changes already made but not committed are rolled back, and all committed transactions, which may not have been physically written to

disk, are redone. These actions ensure the integrity of the database. The active logs reside in the database log path directory.

Roll-forward recovery uses logs in an additional way to allow a database to be rebuilt to a specified point in time. In addition to using the information in the active logs to rebuild a database, the roll-forward function uses archived logs to reapply previous changes.

DB2/2 Version 1 Backup Utilities

Backup/Recover
DB2/2 Version 1.1 provides a utility called *backup* for offline backup. Two types of backup are provided: a full backup of the entire database and an incremental backup of only the changed files (called *changes only*). For offline backup, no other processes can be connected to the database at the same time. Each application must be logged off. When the backup operation starts, SQL operations are not permitted.

A user exit program can be used in conjunction with the backup utility. It allows DB2/2 to use OS/2 commands to interact with devices (or ADSM) directly. You cannot perform a changes-only backup if you enable the user exit program. The changes-only backup relies on OS/2 system commands to determine which files are changed and therefore should be backed up. If you enable the user exit program, you must choose other ways to save your data, such as ADSM or the copy command.

To use the backup utility, the database must be a local database.

DB2/2 Version 1.2 provides a significant enhancement to the Version 1.1 backup utility. The database administrator can perform a backup of a database with active connections, using a new quiesce option. Thus users, in addition to the database administrator, remain connected to the database. You can think of allowing these types of connections as somewhere in-between an offline and online backup.

The quiesce option ensures that all activity on a database has been brought to a halt before a backup is taken. The backup waits until current transactions complete and prevents any new transactions from starting. Once a quiesced state has been established, a backup can be taken.

Export/Import
DB2/2 provides an export/import utility to move data. It can be used as a supplement to your backup strategy, but it is not a replacement for the backup utility. There is a risk of introducing inconsistencies into your database because no synchronization of data with logs is performed.

However, the DB2/2 version 1 backup utility backs up entire databases only, so use of the export/import utility is the only way to capture individual tables.

DB2/2 Version 2 Enhancements

Table Space Concept
With DB2/2 Version 2 databases can be partitioned into *table spaces.*

Backup/Recovery Enhancements
DB2/2 Version 2 provides significant enhancements to the backup/restore utilities. It adds the backup support already available in DB2/6000 V1, including online backups.

Like DB2/6000 V2, DB2/2 V2 can perform backup and recovery at the table space level. If a table space contains a single table, the backup or recovery is equivalent to a table level backup or recovery. Table spaces can be online or offline during the backup process. During recovery, all table spaces other than the one being recovered can remain online.

Another enhancement to the backup/recovery process is *parallel backup and recovery*. It is possible to perform the backup or recovery of a database or tablespace in parallel to or from multiple devices. This enhancement can drastically reduce the elapsed time requirements.

New Load Utility
The new load utility significantly increases the speed of doing data loads and ensures recoverability of the data being loaded. It is intended for either bulk loading of new tables or appending large amounts of data to existing tables. The load utility is restartable and recoverable. If a failure occurs while loading data, you can continue the load without starting from the beginning. To ensure recoverability, a backup copy of the loaded table is created during the load processing.

DB2/2 Backup Alternatives

Backup Type	Database Mode	Intermediate File Created	Backup Granularity	Platforms Supported
DB2/2 V1.1 Backup	Offline	Yes	Database, Changes Only	OS/2
DB2/2 V1.1 Backup with User Exit Program	Offline	Yes	Database	OS/2
ADSM Directly	Offline	No	Database	OS/2, NT
DB2/2 V1.2 Backup	Offline, Quiesced	Yes	Database, Changes Only	OS/2
DB2/2 V1.2 Backup with User Exit Program	Offline, Quiesced	Yes	Database	OS/2
DB2/2 V2 Backup Using ADSM API	Online, Offline	No	Database, Table Space	OS/2, NT

DB2/2 Backup Alternatives
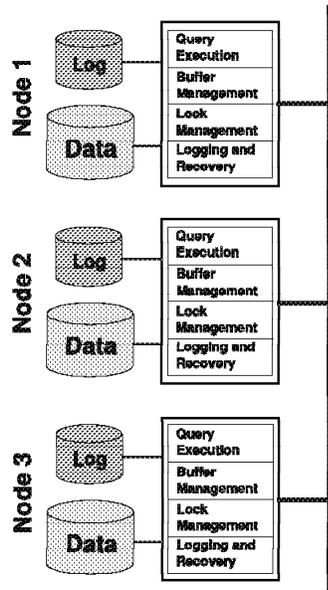
This graphic gives an overview of the DB2/2 backup alternatives with ADSM and summarizes some of their main characteristics, such as:

- Permitted database mode (online or offline)

- Creation of an intermediate OS/2 file (yes or no)

- Backup granularity (the entities that can be backed up)

The best alternative is shown at the bottom of the matrix, DB2/2 V2.1 backup using ADSM API.
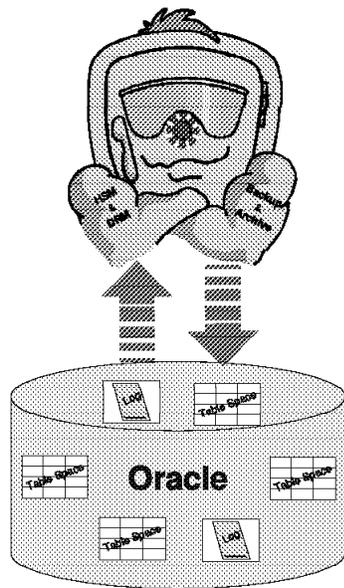
ADSM and Oracle

When we refer to Oracle in this section we mean the Oracle7 Server.

Oracle Structure

The key components of an Oracle database that you must consider for backup and recovery are:

System Tablespace

The system tablespace is automatically created when the database is created. It contains the data dictionary tables for the entire database. The data dictionary is a read-only reference about the database and includes Oracle users′ names, users′ privileges, table names, space usage information, auditing information, and other general database information. The data dictionary is critical to the operation of the database because it records, verifies, and conducts ongoing work. The system tablespace is always online and cannot be taken offline because the data dictionary must always be available to Oracle. We recommend that you reserve the system tablespace for use by Oracle; do not create any user tables in this tablespace.

User Tablespaces

The user tablespaces contain all of the user data. This is where the bulk of the "real" database data is held.

Online Redo Logs

The redo log is a set of files that record all changes made to the database so that in the event of a failure database updates are not lost. The redo log is comprised of two parts: the online redo log and the archived redo log. The online redo log consists of the current log files that are being used to record database changes. Optionally, filled online redo logs can be archived before being reused.

Archived Redo Logs

If you have archived redo logs, the database can be recovered from both instance and disk failures (instead of only instance failures) and backed up while it is open. To use archived redo logs, you run with the ARCHIVELOG mode set to on.

Control Files

Control files, among other things, keep information about the physical structure of the database and log files. If you make a change to the database's physical structure by adding, deleting, or renaming data or log files, you should back up the control files.

Initialization Parameter Files

The initialization parameter files contain instance configuration parameters, such as how much memory to use and what to do with filled online redo logs.

Configuration Files

Configuration files contain options that are common to multiple initialization parameter files.

We recommend that you:

- Use the LVM to mirror at least the disk files used for the system tablespace, control files, and online redo logs.

- Place at least one copy of the system tablespace, control files, and online redo logs on a physical volume separate from the rest of the database.

Oracle Backup Utilities

Alter Tablespace

There are two alternatives for using the alter tablespace utility. Alter tablespace provides a way of taking a tablespace offline while the rest of the database remains in use and a way of backing up a tablespace while it is being updated (online). Forward recovery (using redo logs) is provided on restored tablespaces. You can think of alter tablespace as a way of "freezing" the tablespace such that the database has knowledge of where the backup begins. Forward recovery can provide consistency for any updates that occur after the tablespace is frozen.

To take backups for one or more tablespaces at a time, the database must be operating in ARCHIVELOG mode.

Control File Backup

After making a structural modification (for example, adding a tablespace) to the database, you must back up, at a minimum, the control file. You can use the "alter database backup controlfile" command to back up the control file. The database must be operating in ARCHIVELOG mode.

Export/Import

Oracle provides a pair of utilities for importing and exporting data. You can use these utilities to export and import:

- An entire database

- All of the database objects owned by a particular user

- Individual tables

An incremental option is also provided. After entering the export command and directing the output to a file, you use ADSM commands to back up that file to ADSM storage.

Although you use export while the database is open and available for use, if you use export as a means of backup, all data must be exported in a logically consistent way so that the backup reflects a single point in time. No one should make changes to the database while the export takes place. Ideally, you should run the database in restricted mode during the export so that regular users cannot access the data.

Oracle Enterprise Backup Utility - EBU (also called OBACKUP)

EBU, previously known as the Oracle Parallel Backup utility and sometimes referred to by its main command name OBACKUP, is Oracle's backup utility. EBU is a separate, installable product (available from Oracle) with no additional cost (P/N-A43733-1).

EBU provides backup and restore functions for Oracle7 databases. Using EBU you can perform full and partial, offline and online backups. A full backup consists of all database data files, online redo log files, and Oracle control and parameter files. An offline backup is taken after the database is cleanly shut down. An online backup is taken while users have access to the database; an online backup is also known as a fuzzy, inconsistent, or hot backup. Therefore, with online backups the online redo logs must also be archived.

Whether a backup is online or offline depends on the state of the database at the time of the backup, not the state of the tablespace.

Once you have identified which database you want to back up, EBU locates all of the necessary files and can send them to ADSM.

EBU works in conjunction with media management products such as ADSM. EBU provides an API. Oracle asks each media management vendor to supply the backup media interface. The EBU API consists of call definitions. The media management vendor must write code to implement those calls. The IBM ADSM development team wrote the ADSMConnect Agent for Oracle Backup, which provides an interface between the EBU API calls and the ADSM API routines. The ADSMConnect Agent for Oracle Backup translates your EBU commands into ADSM API calls. The ADSMConnect Agent for Oracle Backup is also known as the *API glue code.*

EBU allows multiple data files or tablespaces to be backed up simultaneously on multiple devices in parallel, thus greatly reducing the time needed to perform backups. Data can also be restored in parallel from multiple devices to ensure speedy restoration.

SQL-BackTrack for Oracle

SQL-BackTrack, from DataTools, Inc., provides a complete database backup solution for Oracle and Sybase databases. It uses the ADSM API to store Oracle and Sybase database objects directly within the ADSM server; an intermediate file is not created. SQL-BackTrack offers facilities such as true incremental backup, object-level backup and recovery, data compression, encryption, and integration directly with ADSM on Sun, AIX, and HP-UX platforms.

DataTools plans to provide SQL-BackTrack for Oracle on Sequent PTX, Siemens-Nixdorf Sinix, DEC Alpha Digital UNIX, and Intel PC Windows NT platforms.

DataTools can be contacted by e-mail at sales@datatools.com or support@datatools.com, or by phone in the United States at (800) 721-8665, (415) 842-9100, or (516) 766-6565.

Oracle Backup Alternatives

| Backup Type	Database Implementation	Database Mode	Intermediate File Created	Backup Granularity	Platforms Supported
dd	Raw Device	Offline	Yes	Raw Device	Any
ADSMPIPE	Raw Device	Offline, (Online with Alter Tablespace)	No	Raw Device	Any
Oracle Export	Raw Device, File System	Online, but Not in Use	Yes	Database, User Objects, Table, Incremental	Any
ADSM Directly	File System	Offline	No	Database, Tablespace, Log files	Any
Oracle Alter Tablespace	File System	Online, Offline	No	Tablespace	Any
EBU with ADSM Connect Agent	File System, Raw Device	Online, Offline	No	Database, Tablespace, Log Files	AIX Sun (Future)
DataTools SQL-BackTrack	File System, Raw Device	Online, Offline	No	Database, Incremental, Tablespace, Data File, Export	AIX, HP, Sun Soon: Digital Unix, NT, Sequent, Sinix

Oracle Backup Alternatives

This graphic gives an overview of the Oracle backup alternatives with ADSM and summarizes some of the main characteristics such as:

- Permitted ways of implementation (on file systems or raw devices)

- Permitted database modes (online or offline)

- Creation of an intermediate file (yes or no)

- Backup granularity (the entities that can be backed up)

Oracle Parallel Server
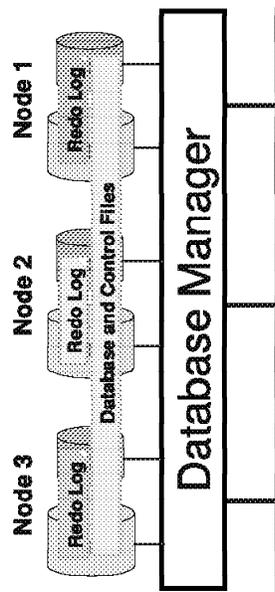
＊ Parallel Database Architecture

◆ **Database Stored across a Network of Processors**
  ‣ Shared Data Files and Control Files
  ‣ Each Oracle Instance Has Its Own Redo Log

◆ **Distributed Lock Manager**

◆ **Not Linearly Scaleable**
  ‣ May Perform Better for Smaller Configurations
  ‣ Not Linearly Scaleable for Larger Configurations

◆ **Backup Tools Must Be Updated to Support Shared Resource Environment**

Oracle Parallel Server

In this section we describe the Oracle Parallel Server.

Parallel Database Architecture

Oracle Parallel Server (OPS) supports two types of loosely coupled architectures: clustered, such as the IBM HACMP/6000, and massively parallel, such as the IBM SP2. OPS is supported on raw devices only.

Database Stored aross a Network of Processors

In an OPS environment, all instances share the database data files and control files but have their own redo log files, known as *threads of redo*. Multiple Oracle instances use Parallel Cache Management to access common data without overwriting each other.

Distributed Lock Manager

Oracle′s Distributed Lock Manager manages the buffers in the various caches. This locking is independent of the Oracle transaction and row-level locking. The combination of the two locking mechanisms allows, for example, two transactions on two nodes to modify different rows in the same data block without contention.

Not Linearly Scaleable

This implementation of parallel database architecture may perform better for small configurations because it is not linearly scaleable for large configurations.

Backup Tools Must Be Updated to Support Shared Resource Environment

You can back up the OPS with ADSM, using SQL-BackTrack on AIX, HP-UX or Sun, or you can use another ADSM raw device backup technique, such as dd, ADSMPIPE (available on any platform), or export. EBU support with ADSM is planned with the ADSMConnect Agent for Oracle Backup.
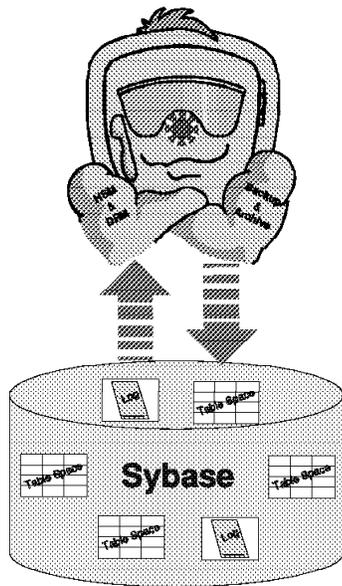
ADSM and Sybase

In this section we describe the Sybase database structure and backup utilities.

Sybase Structure

The key components of a Sybase database that you must consider for backup and recovery are:

The Master Database

The master database controls the user databases and the operation of the server as a whole. It contains system tables that keep track of server information, such as users, databases, storage space, and locks. System tables can also be thought of as the data dictionary or system catalog.

The User Databases

The user databases contain all of the user data as well as a subset of the system tables with information specific to that database.

Other Databases, Such As Model, System Procedure, and Audit

The model database is used as a template for creating user databases. The system procedure database holds system procedures. The audit database is used for security.

Transaction Logs
Each database contains transaction logs. They are really system tables called syslogs. A syslog automatically records every transaction issued by each user of the database.

Segments
Segments are subsets of the database devices available to a particular server database. Each database can have up to 32 segments. Segments provide a flexible tool for assigning objects to particular database devices.

Sybase Backup Utilities
Sybase provides some utilities to backup its databases:

Dump Database
Dump database takes database images of the entire database, including both the data and transaction logs. It allows dynamic dumps; users can continue to use their applications to query or update database data while the dump takes place. In other words, dump database is an online dump utility.

Dump Transaction
Dump transaction copies the transaction log, providing a record of any database changes made since the last database or transaction log dump. It works like an incremental backup by copying all committed transactions since the last dump transaction. So, for example, you might take a full database backup once a week using dump database but back up the transaction logs only during the week with dump transaction. Dump transaction, like dump database, works in online mode.

Load Database and Load Transaction for Recovery
Load database loads the dump into an existing or new database, and load transaction rebuilds the database by reexecuting changes recorded in the transaction log. The load commands are executed while the database is offline.

Bulk Copy Utility
The bulk copy utility is not designed as a backup utility but you can use it to supplement your backup strategy. The bulk copy utility is used to transfer data between the server and the operating system. It is frequently used to copy data into another program, such as another DBMS or spreadsheet. The bulk copy utility is similar to the export and import utilities that other RDBMSs provide.

The bulk copy utility is the only Sybase mechanism to copy data at a table level; the dump utilities work at the database level. You could bulk copy tables to file system files and then use ADSM to back up the file system files. Keep in mind that you must be very careful when using an alternative like this because you run the risk of introducing database inconsistencies.

SQL-BackTrack for Sybase
SQL-BackTrack for Sybase, a product from DataTools, Inc., offers additional database administrative functions and more flexibility to back up Sybase databases.

Physical Format Backups

Physical format backups are dependent on the hardware and software configuration. They can be recovered only to the configuration of the source server.

SQL-BackTrack's physical format backups extend the basic functions of the Sybase dump database command and provide the following options not available from Sybase:

- Incremental backups

  SQL-BackTrack supports the following types of backups: full, true incremental backups, eight levels of differential backups, and transaction log backups.

- Object-level recovery

  SQL-BackTrack lets you recover an individual object (such as a table) from a full, physical backup. This saves considerable time when only a small part of the database is recovered.

Physical format backups provide disaster recovery when used in conjunction with the Sybase transaction log. Roll-forward recovery is provided, so the database can be backed up while it is online.

Logical Format Backups

Logical format backup and recovery capabilities read and write database backups in a format that logically describes the database structure and its data. Thus data can be restored to different hardware or software versions.

Logical format backup and recovery work at the object level. Logical format backups are mainly used for data archiving and data migration. Lock the database from updates while performing logical format backups.

Data Compression and Encryption

By compressing data before writing it, SQL BackTrack saves storage media space and reduces media handling. Optional encryption is the only way of safeguarding sensitive data. Data compression is especially effective for logical format backups (with compressed data taking about 10% of the space of a full physical backup).

Backup Automation

SQL-BackTrack automates the normally interactive Sybase dump procedures. The database administrator simply lists the servers, databases, and/or database objects to back up and provides the correct password in a control file. SQL-BackTrack can then run unattended at any predetermined time.

Media and Backup Tracking

SQL-BackTrack records the time, data, and location of each backup in an online backup history, freeing you from tracking this information manually. In addition, each backup file includes the database allocation information, which otherwise must be tracked manually.

The online backup history makes recovery simple. You simply indicate the database or object to recover, and SQL-BackTrack automatically loads or prompts you for the correct media and applies all incremental or transaction log backups up to a point in time that you specify.

Requirements

SQL-BackTrack supports the following platforms:

- Sun SPARC running SunOS 4.1.x and above or Solaris 2.2 and above

- IBM RS/6000 running AIX 3.2 and above

- Hewlett-Packard 9000 Series running HP/UX 9.0 and above

- DEC Alpha Digital UNIX running 3.2 and above, with optional support of System 11 and above (soon to be supported)

- NCR UNIX 2.0 and 3.0 (soon to be supported)

The DBMS requirements are:

- Sybase SQL Server 4.0.1 or higher

- Sybase Open Client DB-Library

Integrated with ADSM API: ADSM Module

An optional SQL BackTrack module is available from DataTools to integrate SQL BackTrack database backups with ADSM.

DataTools can be contacted by e-mail at sales@datatools.com or support@datatools.com, or by phone in the United States at (800) 721-8665, (415) 842-9100, or (516) 766-6565.

Sybase Backup Alternatives

Backup Type	Database Implementation	Database Mode	Intermediate File Created	Backup Granularity	Platforms Supported
dd	Raw Device	Offline	Yes	Raw Device	Any
ADSMPIPE	Raw Device	Offline	No	Raw Device	Any
Dump Database	Raw Device, File System	Online, Offline	Yes	Database, Transaction Logs	Any
Dump Transaction	Raw device, File System	Online	Yes	Transaction Logs	Any
Bulk Copy Utility	Raw Device,	Online	Yes	Table	Any
ADSM Directly	File System	Offline	No	Database	Any
DataTools SQL-BackTrack Using ADSM API	Raw Device, File System	Online, Offline	No	Database, Incremental, Database Objects	AIX, HP, Sun Soon: Digital UNIX, NCR UNIX

Sybase Backup Alternatives

This graphic gives an overview of the Sybase backup alternatives with ADSM and summarizes some of the main characteristics such as:

- Permitted ways of implementation (on file systems or raw devices)
- Permitted database modes (online or offline)
- Creation of an intermediate file (yes or no)
- Backup granularity (the entities that can be backed up)

ADSM and INFORMIX-OnLine
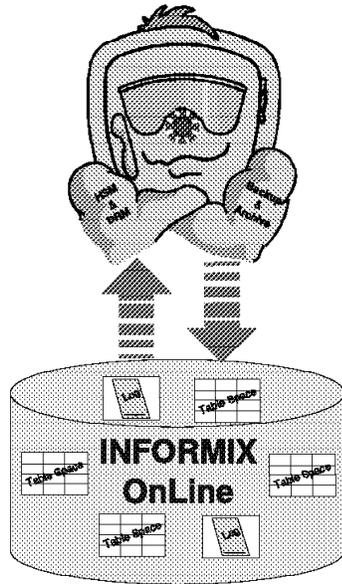
- INFORMIX-OnLine Structure

- INFORMIX-OnLine Backup Utilities
  - Tbtape
  - Ontape
  - Onarchive
  - ON-Bar
  - Tbunload/Tbload (Onunload/Onload)
    and Dbexport/Dbimport

- INFORMIX-OnLine Extended
  Parallel Server (XPS)

- SQL-BackTrack for INFORMIX-OnLine

ADSM and INFORMIX-OnLine

In section we describe the INFORMIX-OnLine database and the methods available to back it up.

INFORMIX-OnLine Structure

The key components of an INFORMIX-OnLine database that you must consider for backup and recovery are:

- The root dbspace

- Noncritical dbspaces

- Logical logs (critical dbspace)

- Physical logs (which are part of the root dbspace by default)

- Configuration files

Root Dbspace

The root dbspace is critical because it contains reserved pages and internal tables that describe, control, and track system information. It can also contain the physical and logical logs, and any user databases or tables you want to store in the root dbspace.

There are other critical dbspaces besides the root dbspace, such as any dbspace that contains the physical and logical logs.

Noncritical Dbspaces

All other dbspaces are called noncritical dbspaces, and they contain all of the user data. This is where the bulk of the "real" database data is held.

Please note that the term *tblspace* for the INFORMIX-OnLine product refers to all of the disk space for a table. This meaning differs from that of tablespace in the other RDBMS products. Dbspace in INFORMIX-OnLine is comparable in meaning to table space in other RDBMS products.

Logical Logs

The logical logs contain a history of the database changes. They are used to recover a database to logical consistency, by rolling forward all committed transactions that have occurred since the checkpoint and rolling back all transactions that were left incomplete. Logical logs also contain changes that are made to the configuration, such as changes to dbspaces and chunks (same meaning as data files in other RDBMSs).

Continuous logical log backups automatically back up each logical log file as it becomes full. We recommend that you have an offline copy of every log except the current one. Should your system fail, you would lose only the modifications held in the current log.

Once a logical log has been backed up, it can be reused provided that no long transactions are held in it and a checkpoint has been done after the logical log backup. Reusing a logical log minimizes the amount of disk space used for logical logs.

You must use the same utility to back up your logical logs as you use to create your database backups. If you use ontape to create your dbspace backups, you must use ontape to back up your logical logs.

Physical Logs

The physical logs are used to recover a database to the most recent point of known physical consistency, the most recent checkpoint. The physical log contains a "before image" of all physical pages that have been modified since the last checkpoint.

The physical logs are also used to enable a backup to take place with the database online. Such a backup will be a snapshot of the database.

The backup utility forces a checkpoint and then starts to back up database pages, all of which are unmodified. If a page is modified during the backup process, an unmodified copy is written to the physical log and the modified copy is marked as "dirty." When the backup program reaches a dirty page, it takes the unmodified copy from the physical log.

Should a checkpoint occur during the backup process, the checkpoint process causes the backup process to scan the physical log for pages that have not been backed up. Once these pages have been backed up, the checkpoint completes, and the backup program continues to back up unmodified pages until the snapshot is complete.

Configuration Files

The configuration files contain all configuration parameters concerning the root dbspace, disk mirroring, physical and logical logs, message files, and tape devices; identification parameters; Shared Memory parameters; and machine-specific parameters.

INFORMIX-OnLine Backup Utilities

Several INFORMIX-OnLine utilities can be used for backup. Below we describe the utilities for the different INFORMIX-OnLine versions.

You can use other INFORMIX-OnLine utilities for backup, but they are actually data migration utilities and not a substitute for tbtape, ontape, or onarchive. The tbunload (or onunload in INFORMIX-OnLine Version 6.0) and dbexport/dbimport utilities are not coordinated with the information stored in logical log files, and they do not save a copy of system information important to INFORMIX-OnLine.

Tbtape

Tbtape is the main backup utility for INFORMIX-OnLine Version 5. It is used to back up logical logs (including continuous backups), change the database logging status, create and archive, or restore data.

Continuous logical log backups automatically back up each logical log as it becomes full.

Database logging is the process of writing transactions to the log files. Different types of logging are supported, such as buffered and unbuffered. Tbtape is used to turn database logging on and off and change the type of logging used.

An archive (in INFORMIX-OnLine terminology) backs up all of the data managed by the database server, which includes *all* dbspaces. This is called a *full system archive*. It is written on a tape or to a file and includes all used disk pages. The archive does not include configuration files, so you have to back them up separately.

Archiving is rather sophisticated (or complicated!) in that it provides for incremental archives so that you do not have to archive the entire database each time. Three levels of archive are provided:

- Level-0, the baseline archive because all unused pages are archived

- Level-1, all changes since the last level-0 archive

- Level-2, all changes since the last level-1 archive

You build a schedule for incremental archives that fits the needs of your environment. For example, you can perform level-0 archives monthly, level-1 archives weekly, and level-2 archives daily, or you can build a more frequent schedule.

The tbtape archive is performed when the database is in online or quiescent mode. Remember that quiescent mode can be thought of as "almost-offline" or a single-user multitasking mode. The tbtape restore is performed when the database is in offline mode.

Ontape

Ontape is one of the main backup utilities in INFORMIX-OnLine Version 6. It is similar to tbtape but provides some significant enhancements. With ontape you can restore specific dbspaces, and the restore can be either warm or cold. With warm restore, INFORMIX-OnLine can be online; with cold restore, INFORMIX-OnLine must be offline. Warm restores can be done for noncritical dbspaces, that is, any dbspace other than the root dbspace and the dbspaces that contain the physical and logical logs. The backup operation includes all dbspaces. Only the restore has changed to allow specific dbspaces.

Another difference between ontape and tbtape is that with ontape you no longer can use the INFORMIX-OnLine onmonitor utility to create the archive. You can execute the ontape utility only from the command line.

Onarchive

Onarchive is a powerful new utility for INFORMIX-OnLine Version 6. It provides the function of ontape but adds significant enhancements including:

- Archive of specific dbspaces and dbspace sets (one or more dbspaces)

- Archive and restore of different dbspaces in parallel

Onarchive provides online and quiescent mode backup and online and offline restore for one or more dbspaces in parallel. You group one or more dbspaces into dbspace sets, and onarchive works against these dbspace sets.

The output produced by ontape and onarchive is incompatible. You cannot create a backup (called an archive in INFORMIX-OnLine) with ontape and restore it with onarchive (or vice versa).

ON-Bar

ON-Bar (Backup and Restore) is new to the INFORMIX-OnLine product line (in Versions 7.21 and 8.0) and enhances the backup and restore facilities. ON-Bar replaces onarchive and ontape and provides the following enhancements:

- Simplifies backup and restore procedures compared to onarchive

  Note that ON-Bar refers to a backup as a backup, not an archive, whereas onarchive refers to a backup as an archive.

- Allows for warm backup and restore of dbspaces

- Enables automatic logical log backup through an ALARMPROGRAM variable that calls a customizable script

- Provides automation and management facilities to reduce operator intervention

- Supports SMP, XMP, and distributed data environments.

ON-Bar is designed to support the X/Open Backup Services API (XBSA). Therefore, any storage manager product that supports this standard works with ON-Bar. The term "storage manager" is used here to denote any type of storage management tool, for example, ADSM. An ADSM X/OPEN API is available for AIX, HP, and Solaris platforms. Thus ON-Bar backups can be sent directly to ADSM through the ADSM X/OPEN API; no intermediate file is created.

Tbunload/Tbload (Onunload/Onload) and Dbexport/Dbimport

INFORMIX-OnLine provides other utilities that can be used as a supplement (not replacement) for the backup utilities.

The tbunload and tbload (onunload/onload) provide the fastest way of moving data, but you cannot modify the database schema (.sql file) or use them to migrate from one hardware platform to another. It is the fastest way because it writes in binary disk-page units. The database schema contains the database definition statements, such as access privileges and object ownerships. The granularity of data supported is either a table or database.

Dbexport and dbimport provide some flexibility; you can modify the database schema, but you must move the entire database.

For both sets of utilities the database server can be online because they obtain exclusive locks on the tables so that no transaction can update the tables at the time of the commands.

INFORMIX-OnLine Extended Parallel Server (XPS)

The INFORMIX-OnLine Extended Parallel Server Version 8.0 (XPS) is a parallel processing version of INFORMIX. It exploits the emerging technology of loosely coupled or shared-nothing computing architectures, including clusters of symmetric multiprocessors (SMPs), for example, the IBM RS/6000 Model J30, and massively parallel processors (MPPs), such as the IBM RS/6000 Scalable POWERparallel Systems (SP2).

XPS is designed to support tasks that require large amounts of data such as data warehousing and decision support. It provides a scalable approach by spreading the processing load across multiple processors or systems and therefore making the best use of available resources.

The tool designed to back up XPS is onbar, which provides an interface between the INFORMIX-OnLine XPS system and a storage manager such as ADSM. Although XPS backup and restore differs from non-XPS backup and restore, INFORMIX has tried to hide the differences from the end user.

SQL-BackTrack for INFORMIX-OnLine

SQL-BackTrack for INFORMIX-OnLine is in beta test at the time of writing.

DataTools can be contacted by e-mail at sales@datatools.com or support@datatools.com, or by phone in the United States at (800) 721-8665, (415) 842-9100, or (516) 766-6565.

INFORMIX-OnLine Backup Alternatives

Backup Type	Database Implementation	Database Mode	Intermediate File Created	Backup Granularity	Platforms Supported
dd	Raw Device	Offline	Yes	Raw Device	Any
ADSMPIPE	Raw Device	Offline	No	Raw Device	Any
INFORMIX Dbexport	Raw Device, File System	Online	Yes	Database	Any
Tbunload/ Onunload	Raw Device, File System	Online	Yes	Database, Table	Any
Onarchive	Raw Device, File System	Online, Quiesced	Yes	Database, Dbspace, Incremental	Any
ADSM Directly	File System	Offline	No	Database	Any
ON-Bar Using ADSM X/Open API	Raw Device, File System	Online, Quiesced	No	Database, Dbspace, Incremental	AIX, HP, Solaris
Future: DataTools SQL-BackTrack	Raw Device, File System	Online, Quiesced	No	Database, Dbspace, Incremental	AIX, HP, Sun

INFORMIX-Online Backup Alternatives

This graphic gives an overview of the INFORMIX-Online backup alternatives with ADSM and summarizes some of the main characteristics such as:

- Permitted ways of implementation (on file systems or raw devices)

- Permitted database modes (online or offline)

- Creation of an intermediate file (yes or no)

- Backup granularity (the entities that can be backed up).

ADSM and SAP R/3
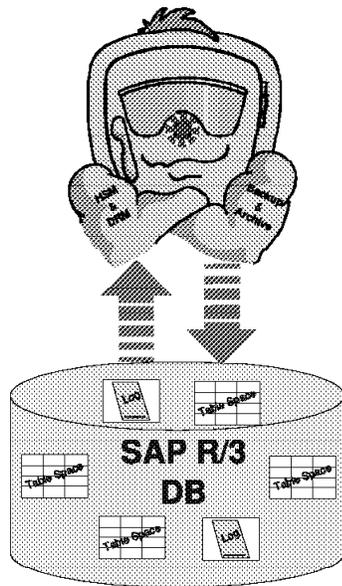
* What Is SAP R/3?

* SAP R/3 Backup Utilities

* SAP R/3 Backup with ADSM
  * BACKINT/ADSM for Oracle
  * ADINT/ADSM for ADABAS

ADSM and SAP R/3

In this section we describe SAP R/3 and the backup utilities available for it.

What Is SAP R/3?

SAP (Systeme, Applikationen und Produkte in der Informationsverarbeitung) is one of the worlds largest software companies providing packaged commercial applications. Founded in 1972 in Germany by five former IBM employees, it now has 28 subsidiaries and affiliates in all of the major industrialized countries of Europe, North America, Asia, and Africa and supports more than 3500 customers in 36 countries worldwide. In July 1993, IBM and SAP formed an alliance by signing an international agreement that covers mutual cooperation in development, marketing, sales, and support of customer business solutions.

SAP R/3 is an integrated client/server package covering accounting, human resources, logistics, and production planning. It also provides an application development environment with a fourth-generation language, ABAP/4.

The SAP R/3 software, the ABAP/4 application programs, and user data are stored on the database server. SAP has chosen a centralized database concept because it is easier to maintain.

SAP R/3 runs on selected platforms from certain hardware vendors. These vendors are called SAP hardware partners. At this time IBM, BULL, DEC, HP, SNI, and Sun are hardware partners of SAP.

The following database products are supported or in beta test with SAP R/3 on AIX/6000:

- DB2/6000 and DB2PE

- Oracle and Oracle Parallel Server

- INFORMIX-OnLine and INFORMIX-OnLine Extended Parallel Server (XPS)

- ADABAS-D (previously Entire-SQL)

SAP R/3 Backup Utilities

SAP R/3 can use DB2, Oracle, INFORMIX-OnLine, and ADABAS databases. You can back up a SAP R/3 database using any backup technique that works for those database products. If the underlying database is an Oracle database, R/3 provides its own integrated database administration utility, SAPDBA, which provides backup and restore functions.

SAP R/3 Backup with ADSM

BACKINT/ADSM is a backup interface utility for backing up SAP R/3 Oracle databases.

ADINT/ADSM is a backup interface utility for backing up SAP R/3 ADABAS-D databases. Both utilities build an interface between SAP R/3 or database commands to ADSM.

SAP R/3 Backup Alternatives

Database	Backup Type	Database Implementation	Database Mode	Intermediate File Created	Platforms Supported
Oracle	BRBACKUP and BRARCHIVE Using BACKINT/ADSM	File System, Raw Device	Offline, Online	No	AIX, Digital UNIX, HP, NCR Unix, Sinix, Solaris Soon: NT
Oracle Parallel Server	Soon: BACKINT/ADSM	Raw Device	Offline, Online	No	
INFORMIX-OnLine	OnArchive Soon: ON-Bar	Raw Device	Offline, Online	Yes	AIX, HP Solaris
DB2/6000 and DB2PE	DB2/6000 Backup Using ADSM API, DB2admin Tool	File System	Offline, Online	No	AIX, HP, NT, Sinix, Solaris
ADABAS-D	ADINT/ADSM	Raw Device	Offline, Online	No	AIX, HP NT, Sinix, Solaris On request: Digital UNIX, NCR UNIX

SAP R/3 Backup Alternatives

There are several methods of backing up databases in a SAP R/3 environment, but we recommend using a method that is officially supported by SAP. At the time of writing, SAP supports the following methods of backing up a database with ADSM:

- For DB2 for AIX and DB2PE, the database utilities (which use the ADSM API) and the SAP R/3 DB2admin tool are supported. The preferred method is the SAP R/3 DB2admin tool.

- For INFORMIX-OnLine, only the *ontape* and *onarchive* database utilities are supported. Support is planned for the new ON-Bar utility for INFORMIX-OnLine and INFORMIX-OnLine XPS. ON-Bar uses the ADSM X/OPEN API.

- For ADABAS-D, the ADABAS-D to ADSM backup interface, *ADINT/ADSM*, is supported. ADINT/ADSM, a service offering, uses the ADSM API.

- For Oracle, the SAP R/3 Oracle to ADSM backup interface, *BACKINT/ADSM*, is supported for Oracle databases installed on file systems, raw devices, and the OPS. BACKINT/ADSM V2.1.1.0 uses the ADSM API.

ADSM and Lotus Notes
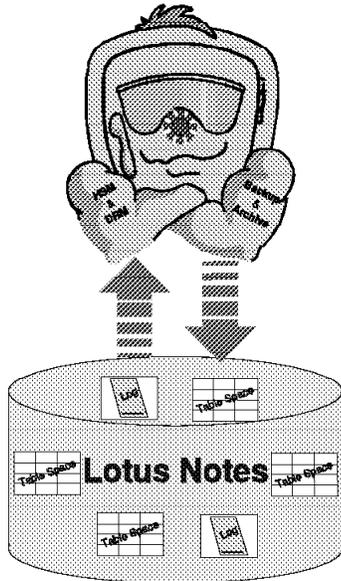
- ✴ **What Is Lotus Notes?**

- ✴ **ADSM Backup/Archive Client with Lotus Notes**
  - • Entire Database Backup Only
  - • Database Contained in Single Flat File
  - • Database Can Be Online but Must Not Be In Use

- ✴ **ADSM Lotus Notes Backup Agent**
  - • Incremental Backup with Note and Document Granularity
  - • Restore
    - ‣ Restore Using Lotus Notes Workspace
    - ‣ Command Line for Backups and Restores

ADSM and Lotus Notes

In this section we describe how to use ADSM to back up Lotus Notes. A detailed redbook on the subject is available: *Using ADSM to Back Up Lotus Notes* (SG24-4534).

What Is Lotus Notes?

In the 1960s and 1970s, before the widespread acceptance of PCs, most computers were multiuser minis or mainframes that supported groups of users, enabling those users to share data, programs, and other information. As PCs became popular in the 1980s, users found themselves "doing their own thing" on their own machine, electronically isolated from their coworkers. Sharing information was cumbersome. PCs made individuals productive but did little to further the interests of groups of users. Now software has come full circle and again is enabling groups of people to work together.

- • **Groupware for sharing information**

  Lotus Notes is an example of a growing genre of software called *groupware*—software designed to be used by groups of people for sharing information and working together.

  One way in which Notes enables you to share information is by providing a kind of information repository that you can share with other people. We call this repository a Lotus Notes database or NSF database.

  An NSF database contains a collection of *documents*. A document is the basic unit of information in Lotus Notes. It may be a news article, a description of a customer complaint, a status report, or a telephone message. Notes documents can contain text, graphics, pictures, and other

kinds of information. Some documents are highly structured and may have a very specific format; others are more free-form.

Actually, an NSF database contains different kinds of *notes*. *Documents*, also called *data notes*, is a subclass of *notes* and the main class of notes with which the end user deals. But there are also other classes of notes, such as *form notes* and *view notes*.

- **Enterprisewide client/server solution**

  Lotus Notes is a client/server solution designed for a distributed environment.

  The server is the machine on which shared databases are located and which manages the documents. There can be multiple servers in a Lotus Notes installation holding copies of the same databases at different locations. These databases are synchronized during a process called *replication*, whereby Lotus Notes distributes and updates copies of the same database on different Notes servers.

  The clients that want to share information can be spread across the globe as long as they are connected to a Notes server by a network. Multiple network protocols are supported (LAN, WAN, and ASYNC). Lotus Notes clients can also have their own local Lotus Notes databases, which are not shared with other Lotus Notes clients.

  The Notes server runs on many platforms, such as OS/2, Windows, NetWare, and various UNIX platforms. The Notes client runs on several platforms, such as OS/2, Windows, Mac, and various UNIX platforms. A Notes server is basically the same code as a Notes client, but the client has some of the server code disabled.

ADSM Backup/Archive Client with Lotus Notes

Entire Database Backup Only
You can back up a Lotus Notes database with the ADSM backup/archive client. ADSM directly backs up the entire Notes database.

Database Contained in Single Flat File
The database of Lotus Notes is contained in a single flat file.

Database Can Be Online but Must Not Be In Use
If the ADSM administrator sets up the management classes for NSF databases so that the serialization is set to *shared static*, ADSM will try a number of times to take a valid backup copy of the databases. A fuzzy copy of the databases will not be taken in this situation, and the Notes server does not have to be taken down for the backup.

Restoration of a full database with the ADSM client works best if you use the recommendations for backup and restore in *Lotus Notes Administrator's Guide: Server for NetWare, OS/2, UNIX*. The administrator's guide discusses special considerations if a database is replicated.

ADSM Lotus Notes Backup Agent

Multiple notes can be in a Notes database, and a database can grow so large that it becomes necessary to back up and restore only parts of the database.

The ADSM Lotus Notes Backup Agent (packaged with the ADSM V2 OS/2 client), provides additional function and flexibility over the full Notes database backup. The Lotus Notes Backup Agent exploits the ADSM API and currently runs on OS/2 Notes platforms only. It can run on an OS/2 Notes server or client but can back up only local databases on the platform on which it is running. Support for Lotus Notes databases on Windows NT and AIX is a high priority requirement.

The ADSM Lotus Notes Backup Agent works with either ADSM V1 or V2 servers and uses either the ADSM V1 Driver 7 level of the API or the ADSM V2 API (which comes packaged with the OS/2 backup/archive client).

The ADSM Lotus Notes Backup Agent now supports Notes 4.1. The support is packaged in PTF 4 for OS/2, which is available on the anonymous ftp server. The PTF number is IP20801. The customer should pick up the entire OS/2 client package from the directory. As usual, the directory is /adsm/fixes/v2r1/os2/english for the English version of the OS/2 client. Each of the different OS/2 directories (for different languages) also contains the Lotus Notes Backup Agent image file.

Incremental Backup with Note and Document Granularity

Incremental backups of databases are provided for both nondata and data notes. The first time the incremental backup is performed, in essence a "full" backup is done because every note is new or changed. Nondata notes include document design templates or forms. The Notes database can be online during the backup.

Restore

Each note is treated as a separate ADSM object. Users can restore single documents; an entire database (by rebuilding it from backup copies of notes); a range of notes; and notes that they deleted but are not yet purged from the Notes database.

Restore Using Lotus Notes Workspace: GUI support is provided through the Lotus Notes workspace interface. Users can use the GUI to restore single documents; an entire database (by rebuilding it from backup copies of notes); a range of notes; and notes that they deleted but are not yet purged from the Notes database. Restore of a single document or note can be done only through the GUI.

Command Line for Backups and Restores: There are three new backup and restore commands:

- **dsmnotes incr**

  The incremental backup command

- **dsmnotes restore**

  Use when use of the Lotus Notes workspace interface is inappropriate, for example, when you want to schedule the restore operation. The command line restore can rebuild copies of the database or restore a range of notes to a database file that was itself restored from an ADSM client. Restoration of

single Lotus documents must be done through the Lotus Notes workspace (GUI).

- **dsmnotes restdel**

  Restores backup copies of notes that the user deleted but are not yet purged from the Notes database. These deleted notes can be restored to the original database without restoring the entire database.

Lotus Notes has a facility through which executables can be scheduled. It allows scheduling of dsmnotes incr commands. The Lotus Notes scheduling facility is an alternative to using the new ADSM V2 scheduling function.

Appendix A. Server Details
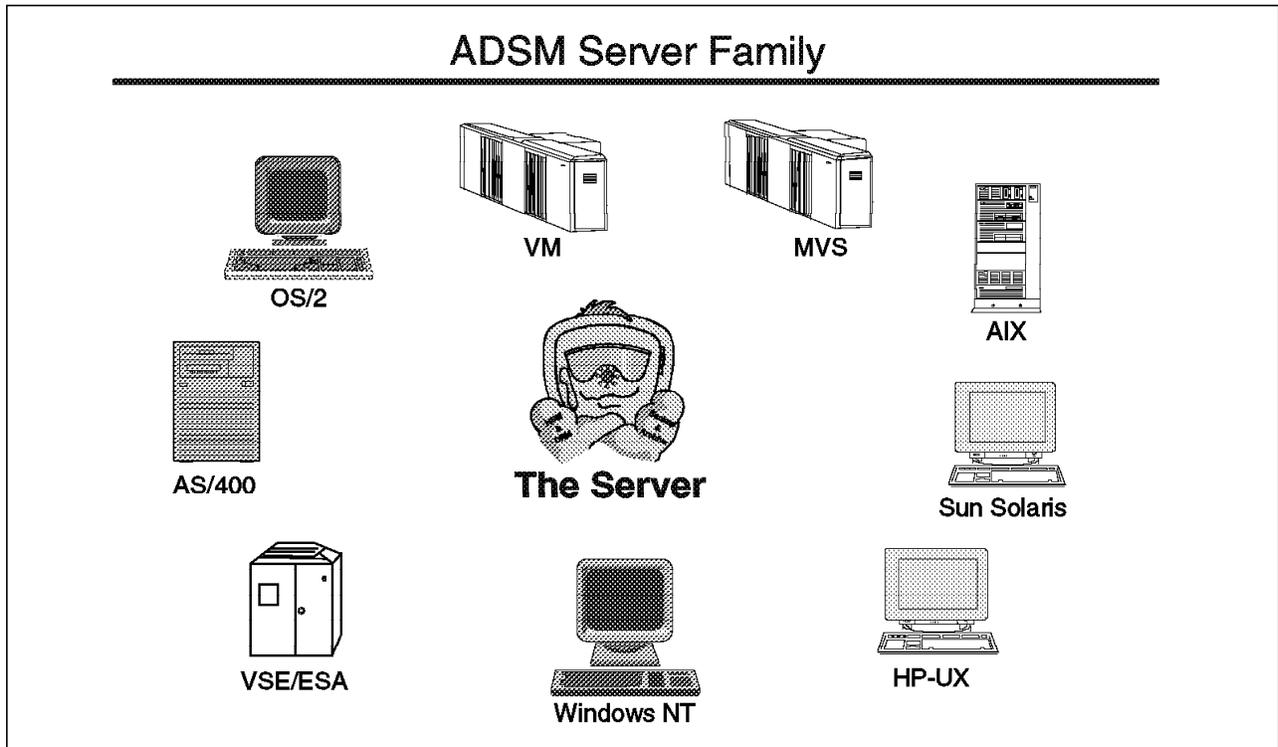
ADSM Server Family

ADSM Server Family

The ADSM server family consists of the following nine servers:

- ADSM for AIX, Version 2.1
- ADSM for AS/400, Version 2.1
- ADSM for HP-UX, Version 2.6
- ADSM for MVS, Version 2.1
- ADSM for OS/2, Version 2.1
- ADSM for Sun Solaris, Version 2.6
- ADSM for VM, Version 2.1
- ADSM for VSE/ESA, Version 1.2
- ADSM for Windows NT, Version 2.1

For some of the servers, the ADSM functions they support depends on the server operating system level.

With these servers, IBM provides customers with a wide range of choices. The ADSM server can be installed on both IBM and non-IBM platforms and can serve the backup-archive requirements for a wide range of clients.

For API clients, some of the ADSM servers have an enhancement known as Partial Object Retrieve (POR). POR gives the application the ability to batch like objects together into a larger superobject and selectively retrieve all or part of the superobject. This ability to retrieve part of an object improves retrieval performance and enables a reduction in storage requirements. This can be very important for data-intensive applications.

The POR function is available for the AIX, AS/400, MVS, OS/2, and VM servers.

As the twenty-first century approaches, you can rest assured that ADSM will support dates in the year 2000 and beyond.

The graphics that follow show the environments in which the ADSM servers run, as well as the clients and some of the devices they support.

For the latest information about the ADSM server family, look up the ADSM home page on the Web, or go directly to the ADSM family page:

- http://www.storage.ibm.com/adsm

- http://www.storage.ibm.com/storage/software/adsm/ads0fam.htm

You can also refer to *ADSM Device Configuration* (SH26-4044) for detailed device descriptions and worksheets. The concept and use of hierarchical storage are described in the redbook entitled *Using ADSM Hierarchical Storage Manager* (SG24-4631). See also the redbook entitled *ADSM API Examples for OS/2 and Windows* (SG24-2588) for further details about API programming.

```
┌─────────────────────────────────────────────────────────────────┐
│                    ADSM for AIX Server                            │
│  ▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬▬   │
│                          Overview                                 │
│                                                                   │
│           ★ Operating System                                      │
│           – ADSM Version 2.1.0    AIX 3.2.5, 4.1 or later         │
│           – ADSM Version 2.1.5    AIX 4.1.4 or later              │
│           ★ Processor Support                                     │
│           – Any RS/6000 System                                    │
│           ★ Communication Protocols                               │
│           – TCP/IP                                                │
│           – APPC                                                  │
│           – NetWare  IPX/SPX                                      │
│           – NetBIOS                                               │
│           – Shared Memory                                         │
│           ★ Functions Supported                                   │
│           – HSM                                                   │
│           – DRM                                                   │
│           – POR                                                   │
│           ★ Product Number                                        │
│           – 5801-AAR  (USA and Canada)                            │
│           – 5639-B93  (Others)                                    │
│           – Modification 2.1.5 for SMP                            │
│                                                                   │
└─────────────────────────────────────────────────────────────────┘
```

ADSM for AIX Server: Overview

The ADSM for AIX server is an ADSM Version 2.1 server. It is a major
development platform for ADSM and supports all ADSM functions.

Hardware features such as LANs and disk drive adapters are transparent to the
ADSM server as long as they are supported by the base operating system.

For the latest information about the AIX server product, requirements, or clients
supported, see the ADSM Web page:

• http://www.storage.ibm.com/storage/software/adsm/ad6serv.htm

Operating System
There are two versions of the ADSM server for AIX. ADSM Version 2.1.0
supports AIX 3.2.5 and AIX 4.1 or later for single processor systems. ADSM
Version 2.1.5 is enabled for SMP systems and supports AIX 4.1.4 and later.

Processor Support
ADSM for AIX runs on any RS/6000 system with at least 32 MB of memory.

Communication Protocols
The ADSM for AIX server supports TCP/IP, APPC, IPX/SPX, and NetBIOS for
communicating with its ADSM clients. For a client and server on the same
system, the Shared Memory communication protocol can be used for better
performance. The communication protocols are transparent to the physical
network.

The following communication protocol levels are required:

• TCP/IP - Support is standard with all versions of AIX.

- APPC (LU6.2)

  - For AIX 3.2.5, support is through IBM AIX SNA Server/6000 2.1 or later.

  - For AIX 4.1 and later, support is through IBM AIX SNA Server/6000 3.1 or later.

- IPX/SPX

  - For AIX 3.2.5, support is through IBM AIX NetBIOS and IPX/SPX Support/6000, or through IBM NetWare for AIX for RS/6000 from IBM V3.11.

  - For AIX 4.1 and later, support is through IBM NetWare for AIX 3.11b through IBM AIX NetBIOS and IPX/SPX support for AIX (AIX 4.1.3 required) or through IBM AIX Connections (AIX 4.1.4 required).

- NetBIOS

  - For AIX 3.2.5, support is through IBM NetBIOS on Token-Ring/6000 or IBM AIX NetBIOS and IPX/SPX Support/6000.

  - For AIX 4.1.3 and later, support is through IBM AIX NetBIOS and IPX/SPX support for AIX or IBM AIX Connections (AIX 4.1.4 required).

- Shared Memory

  Shared Memory communication initializes through the TCP/IP support included in AIX. It is available for a client and server on the same system.

Functions Supported

The ADSM for AIX server supports HSM for these client platforms:

- IBM AIX

- Sun Solaris 2.5

- Novell NetWare

- Windows NT

HSM clients for AIX and Solaris require the Space Management feature to be installed. The AvailHSM (formerly NetSpace) product is required on NetWare, and the OPEN/stor product is required on Windows NT. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

The ADSM for AIX server supports POR.

Product Number

In the USA and Canada, the ADSM for AIX server is available as product number 5801-AAR for one user. For multiple users, the user authorizations product number 5807-AAR must also be ordered. For further ordering information, look up the current announcement letter. Note that although AIX, OS/2, and Windows NT have the same order number, they have different feature numbers.

Outside the USA and Canada, the order number for the AIX server is 5639-B93.

For SMP systems the ADSM Version 2.1.5 feature should be ordered.
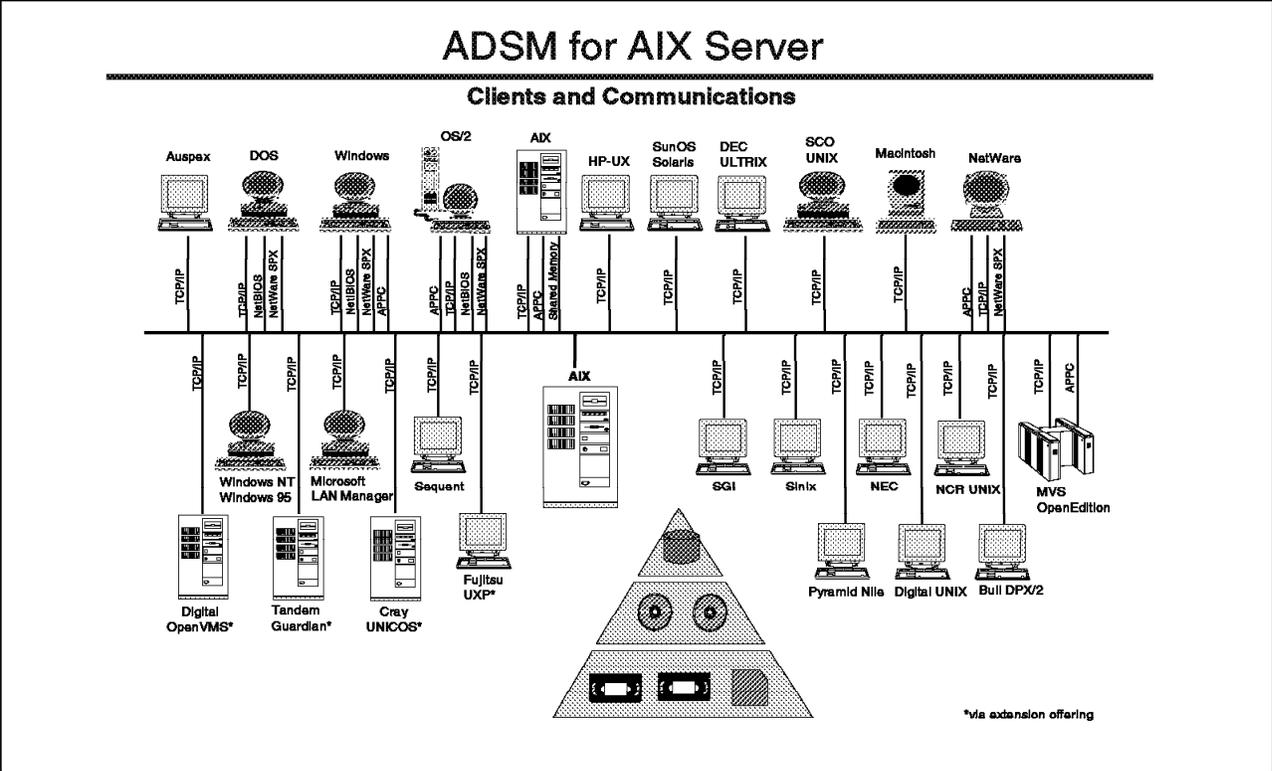
ADSM for AIX Server: Clients and Communications

This graphic shows the backup-archive clients and communication protocols that the ADSM for AIX server supports. Some of these clients use the same client software (for example, Auspex is supported through the Sun Solaris client), and others are supported by extension offerings from IBM or other sources.

The flyer entitled *ADSM for AIX Spec Sheet* (G229-1268) shows the latest clients and protocols that the ADSM for AIX server supports. You can also get the most up-to-date information from the ADSM home page on the Web.

The different UNIX clients require the UNIX Environment Support feature, and the desktop clients require the Desktop Environment Support feature. The UNIX environment is not required for the AIX client on this server. The OpenEdition MVS client requires the Open Systems Environment Support feature.

ADSM for AIX Server

Device Support
such as, but not limited to:

Any Disk Drives Supported by AIX
- Journaled File System
- Raw Logical Volume I/O

4 mm Tape Drives
- IBM 7206
- HP Jetstore x000e
- Sony SDTx000
- WangDat 3x00DX

8 mm Tape Drives
- IBM 7208
- Exabyte EXB
- Dynatek HSB

1/2" Tape Drives
- IBM 3490E
- IBM 3590

DLT Drives
- DEC TZ87, TZ88
- Quantum DLT

Optical Drives
- IBM 0632
- IBM 7209
- Philips LMS

4 mm Tape Libraries
- HP C15xx
- StorageTek 9704
- IBM 7332-005
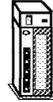
8 mm Tape Libraries
- StorageTek 9708
- Exabyte EXB
- LAGO LS/380L DataWheel
- IBM 7331-205

1/2" Tape Libraries
- IBM 3494 L10
- IBM 3495

DLT Libraries
- ATL products
- Quantum DLT

Optical Libraries
- IBM 3995 Dataserver
- HP 20T
- Philips LMS 6602

ADSM for AIX Server: Device Support

This graphic shows examples of devices that the ADSM for AIX server supports. Most of the devices attach through a standard SCSI, but some of the tape systems also use parallel- or ESCON-channel attachments

This is not a complete list, and new products are continuously being added. For the most up-to-date list of supported devices, call the IBM Information Support Center in the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local IBM representative or reseller. You can also get the most up-to-date information from the ADSM home page on the Web. The flyer entitled *ADSM for AIX Supported Storage Devices* (G225-6620) is continuously updated and shows the latest supported devices.

```
┌─────────────────────────────────────────────────────────────────────────────┐
│                                                                               │
│                       ADSM for AS/400 Server                                  │
│         ═══════════════════════════════════════════════════════════          │
│                            Overview                                           │
│                                                                               │
│                   ★ Operating System                                          │
│                   ─ OS/400  Version 3 Release 2                                │
│                   ─ OS/400  Version 3 Release 7                                │
│                                                                               │
│                   ★ Processor Support                                         │
│                   ─ Any AS/400 System                                         │
│                                                                               │
│                   ★ Communication Protocols                                   │
│                   ─ TCP/IP                                                     │
│                   ─ APPC                                                       │
│                   ─ NetWare  IPX/SPX                                           │
│                                                                               │
│                   ★ Functions Supported                                       │
│                   ─ HSM                                                        │
│                   ─ POR                                                        │
│                                                                               │
│                   ★ Product Number                                            │
│                   ─ 5763-SV2  -  V3.2                                          │
│                   ─ 5716-SV2  -  V3.7                                          │
│                                                                               │
└─────────────────────────────────────────────────────────────────────────────┘
```

ADSM for AS/400 Server: Overview

The ADSM for AS/400 server is an ADSM Version 2.1 server. The ADSM Version
1.2 server is still available for customers using older versions of OS/400.

For the latest information about the ADSM for AS/400 server requirements or
clients supported, see the ADSM Web page at:

- http://www.storage.ibm.com/storage/software/adsm/ad4serv.htm

For more information about installing and implementing the ADSM for AS/400
server, see the redbook entitled *Setting Up and Implementing ADSM/400*
(GG24-4460).

Operating System
The ADSM for AS/400 V2.1 server runs on OS/400 Version 3 releases 2 and 7.
Customers using OS/400 Version 2.3, Version 3.1, or Version 3.6 can still use the
ADSM for AS/400 Version 1.2 server.

Processor Support
The ADSM for AS/400 server runs on any processor that supports the OS/400
operating system.

Communication Protocols
The ADSM for AS/400 server supports TCP/IP, APPC, and IPX/SPX
communication protocols, all of which are included with the current release of
OS/400. For customers still running on OS/400 Version 2.3, only APPC and
TCP/IP together with TCP/IP Connectivity Utilities/400 (5738-TCI) are supported.

Functions Supported

The ADSM for AS/400 server supports HSM for these client platforms:

- IBM AIX
- Sun Solaris 2.5
- Novell NetWare
- Windows NT

HSM clients for AIX and Solaris require the Space Management feature to be installed. The AvailHSM (formerly NetSpace) product is required on NetWare, and the OPEN/stor product is required on Windows NT. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

The ADSM for AS/400 server supports POR.

Product Number

The ADSM for AS/400 server is available as product number 5763-SV2 for OS/400 V3.2 and as product number 5716-SV2 for OS/400 V3.7.
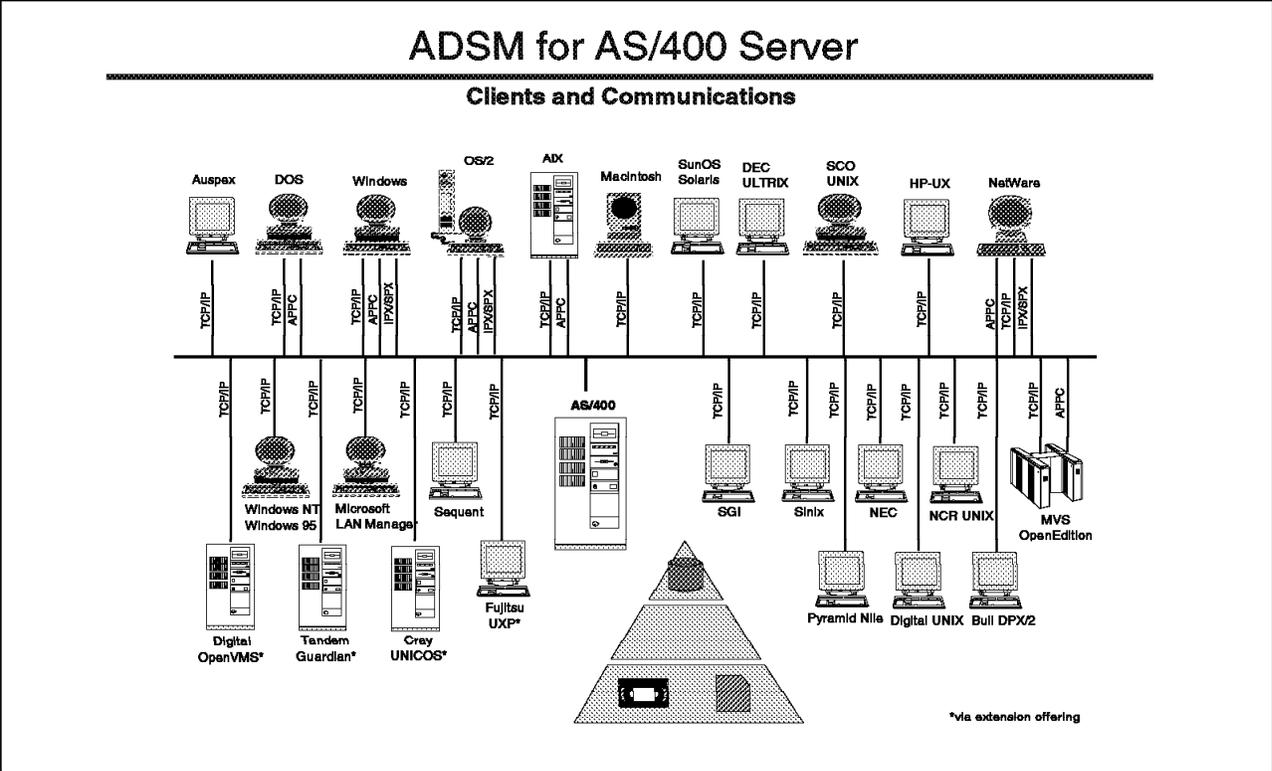
ADSM for AS/400 Server

Clients and Communications

ADSM for AS/400 Server: Clients and Communications

This graphic shows the backup-archive clients and communication protocols that the ADSM for AS/400 server supports.

The DOS and Windows clients are enhanced to support APPC. For Windows you must use Networking Service/Windows (NS/Windows) 1.0 or IBM Client Access/400 for Windows. For DOS use either IBM PC Support/400 for DOS 2.3 together with IBM LAN Support program 1.2.6 or IBM Client Access/400 for Extended DOS 3.0 or later.

The TCP/IP support of the clients is the same as for any other ADSM server.

The flyer entitled *ADSM/400 Spec Sheet* (G229-1270) shows the latest clients and protocols that the ADSM for AS/400 server supports. You can also get the most up-to-date information from the ADSM home page on the Web.

The different UNIX clients require the UNIX Environment Support feature, and the desktop clients require the Desktop Environment Support feature. The OpenEdition MVS client requires the Open Systems Environment Support feature.

ADSM for AS/400 Server

Device Support

such as, but not limited to:

Any Disk Drives Supported by OS/400
– AS/400 Standard DASD
– IBM 9332
– IBM 9335
– IBM 9446
– IBM 9337

8 mm Tape Drives
– IBM 6390 Model 1 (OS/400 V3 only)
– IBM 7208

1/2" Reel Drives
– IBM 2440, 3422, 3430, 9347, 9348

1/2" Cartridge Drives
– IBM 3480, 3490
– IBM 3590 Model B11 (OS/400 V3 only)
– IBM Magstar MP (OS/400 V3.7 only)

QIC Tape Drives
– IBM Model 634x
– IBM Model 636x, 9346
– IBM Model 637x, 6380
   (OS/400 V3 only)

8 mm Tape Libraries
– IBM 9427 Model 210
   (OS/400 V3 only)

Tape Libraries
– IBM 3494 Model L10

ADSM for AS/400 Server: Device Support

This graphic shows examples of devices that the ADSM for AS/400 server supports, including any disk storage devices and tape devices supported by OS/400.

This is not a complete list, and new products are continuously being added. For the most up-to-date list of supported devices, call the IBM Information Support Center in the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local IBM representative or reseller. You can also get the most up-to-date information from the ADSM home page on the Web. The ADSM for AS/400 does not support optical devices. Please note that not all devices are supported by all releases of OS/400, specifically certain models of the IBM 63xx and 720x tape drives and the IBM 3590 and Magstar MP tape cartridge libraries.

Refer to the flyer entitled *ADSM/400 2.3 Device Validation List* (G225-6625) for the devices supported for Version 2.3 of OS/400 and the flyer entitled *ADSM/400 Supported Storage Devices* (G225-6622) for the other versions of AS/400.

```
ADSM for HP-UX Server

Overview


★ Operating System
  – HP-UX level 9.04


★ Processor Support
  – HP9000 series 800 or compatible workstation


★ Communication Protocol
  – TCP/IP


★ Functions Supported
  – HSM
  – DRM
  – POR
   Product Number
  – 5801-AAR (US,. AP and Canada)
  – 5639-B21 (others)
```

ADSM for HP-UX Server: Overview

The ADSM for HP-UX server is an ADSM Version 2.1 server.

For the latest information about the ADSM for HP-UX server requirements or clients supported, see the ADSM Web page at:

- http://www.storage.ibm.com/storage/software/adsm/adhserv.htm

Operating System
The ADSM for HP-UX server runs on an HP-UX level 10.10 or 10.20 operating system and must be installed on a file system. There is no raw logical volume support in the base HP-UX product.

Processor Support
To run the ADSM for HP-UX server, you must use an HP9000 series 800 or compatible workstation.

Communication Protocol
The ADSM for HP-UX server supports TCP/IP, which comes standard with HP-UX.

Functions Supported
The ADSM for HP-UX server supports HSM for these client platforms:

- IBM AIX

- Sun Solaris 2.5

- Novell NetWare

- Windows NT

HSM clients for AIX and Solaris require the Space Management feature to be installed. The AvailHSM (formerly NetSpace) product is required on NetWare, and the OPEN/stor product is required on Windows NT. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

The ADSM for HP-UX server supports POR.

Product Number

In the USA, Canada and Asia Pacific, the ADSM for HP-UX server is available as product number 5801-AAR for one user. For multiple users, the user authorizations product number 5639-B21 must also be ordered. For further ordering information, look up the current announcement letter. Note that although AIX, OS/2, and Windows NT have the same order number, they have different feature numbers.

Outside the USA and Canada, the order number for the HP-UX server is 5639-B21.
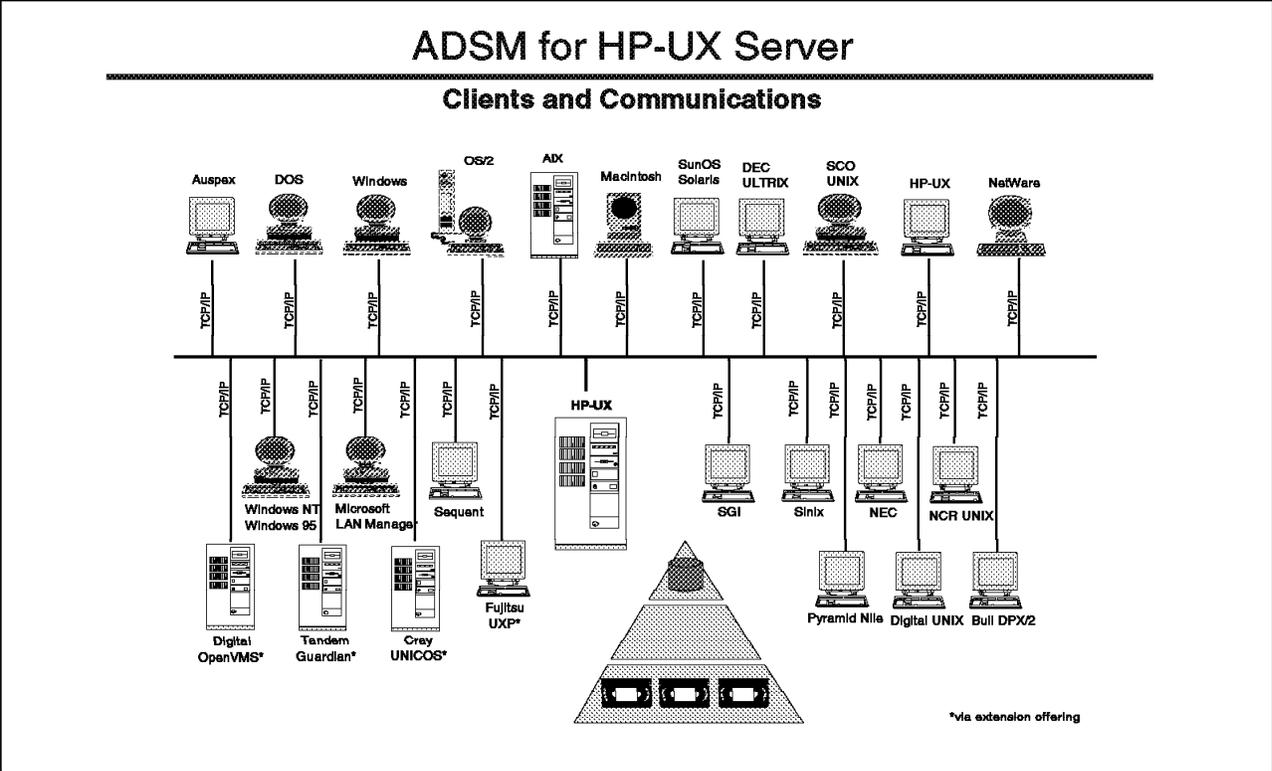
ADSM for HP-UX Server

Clients and Communications

ADSM for HP-UX Server: Clients and Communications

This graphic shows the backup-archive clients and communication protocol that the ADSM for HP-UX server supports.

The flyer entitled *ADSM for HP-UX Fact Sheet* (G229-1275) shows the latest clients supported by ADSM for HP-UX. You can also get the most up-to-date information from the ADSM home page on the Web.

The different UNIX clients require the UNIX Environment Support feature, and the desktop clients require the Desktop Environment Support feature.

ADSM for HP-UX Server

Device Support

such as, but not limited to:

Any Disk Supported by HP-UX 9.04

4 mm Tape Drives
- IBM7206 Model 005
- HP 35470a, 35480a, C1533e
- HP Jetstore 2000e, 5000e, 6000e
- WangDat 3300dx, DDS-2
- WangDat 3400dx, DDS-2
- Sony SDT-5000, DDS-2

8 mm Tape Drives
- Andataco Encore 8205, 8505
- Dynatek HSB-2300, HSB-10.0, HSB-5000
- Exabyte EXB-8200, EXB-8205, EXB-8205XL
- Exabyte EXB-8500, EXB-8500C, EXB-8505, EXB-8505XL
- Transitional Technologies, Inc., CTS-8000H, CTS-8519H

4 mm Tape Libraries
- Diverse Logistics Libra 8, Libra 16
- HP C1553A, C1561A
- HP SureStore 12000e
- IBM 7332-005
- ADIC VLS

8 mm Tape Libraries
- Andataco Encore 10e, 210, 440, 120, 480
- Boxhill Bread Box
- Boxhill Light Box
- Boxhill Freezer Box
- Boxhill Ice Box
- Exabyte EXB-10e, EXB-10l, EXB-60, EXB-210, EXB-440, EXB-120, EXB-480
- Lago LS/380L DataWheel
- Transitional Technologies, Inc., Series 10
- IBM 7331-205

ADSM for HP-UX Server: Device Support

This graphic shows examples of devices that the ADSM for HP-UX server supports, including any disk storage device supported by HP-UX 9.04. It is not a complete list, and new products are continuously being added. For the most up-to-date list of supported devices, call the IBM Information Support Center in the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local IBM representative or reseller. You can also get the most up-to-date information from the ADSM home page on the Web.

The flyer entitled *ADSM for HP-UX Device Validation List* (G225-6627) is continuously updated and shows the latest supported devices. ADSM for HP-UX does not support optical devices.

```
+----------------------------------------------------------------------+
|                      ADSM for MVS Server                             |
|  ==================================================================  |
|                           Overview                                   |
|                                                                      |
|    ★ Operating System            ★ Functions Supported               |
|    – OS/390 Release 1             – HSM                               |
|    – MVS/ESA SP V5.1              – DRM                               |
|    – MVS/ESA SP V4                – POR                               |
|    – MVS/ESA SP V3.1.3                                                |
|                                  ★ Product Number                    |
|    ★ Processor Support           – 5655-119                          |
|    – Any S/390 Architecture Processor                                |
|                                                                      |
|    ★ Communication Protocols                                         |
|    – APPC                                                             |
|    – 3270 Ext                                                         |
|    – TCP/IP                                                           |
|    – IUCV                                                             |
|                                                                      |
+----------------------------------------------------------------------+
```

ADSM for MVS Server: Overview

The ADSM for MVS server is an ADSM Version 2.1 server.

For the latest information about the ADSM for MVS server requirements or clients supported, see the ADSM Web page at:

- http://www.storage.ibm.com/storage/software/adsm/ad1serv.htm

Operating System

The ADSM for MVS server runs on one of the following operating systems:

- IBM OS/390 Release 1 or later

- IBM MVS/ESA SP Version 5 Release 1 or later

- IBM MVS/ESA SP Version 4

- IBM MVS/ESA SP Version 3 Release 1.3

One of the following language run-time environments is a prerequisite for the ADSM for MVS server:

- IBM SAA AD/Cycle Language Environment/370

- IBM C/370 Library Version 2

- MVS C/C++ Language Support feature on MVS/ESA SP Version 5 Release 2 with upgrade for IBM C/C++ for MVS Version 3 Release 1 support

The OS/390 system includes the appropriate language environment, so no additional products are required.

Processor Support

The ADSM for MVS server runs on any S/390 architecture processor.

Communication Protocols

The ADSM for MVS server supports the following communication protocols:

- APPC (LU6.2)

- 3270 Extended Data Stream (LU2.0)

- TCP/IP, available through either IBM TCP/IP 2.1 or later for MVS or Interlink ICS SNS/TCPaccess Version 2.1

- IUCV, available through IBM TCP/IP 2.2 or later for MVS

Functions Supported

The ADSM for MVS server supports HSM for the following client platforms:

- IBM AIX

- Sun Solaris 2.5

- Novell NetWare

- Windows NT

  HSM clients for AIX and Solaris require the Space Management feature to be installed. The AvailHSM (formerly NetSpace) product is required on NetWare, and the OPEN/stor product is required on Windows NT. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

The ADSM for MVS server supports the Disaster Recovery Manager feature, which is described in detail in Chapter 2, "Disaster Recovery Manager" on page 73. Also see the redbook entitled *ADSM for MVS: Recovery and Disaster Recovery* (SG24-4537) for details on how to use DRM.

The ADSM for MVS server supports POR.

Product Number

The ADSM for MVS server is available as product number 5655-119.
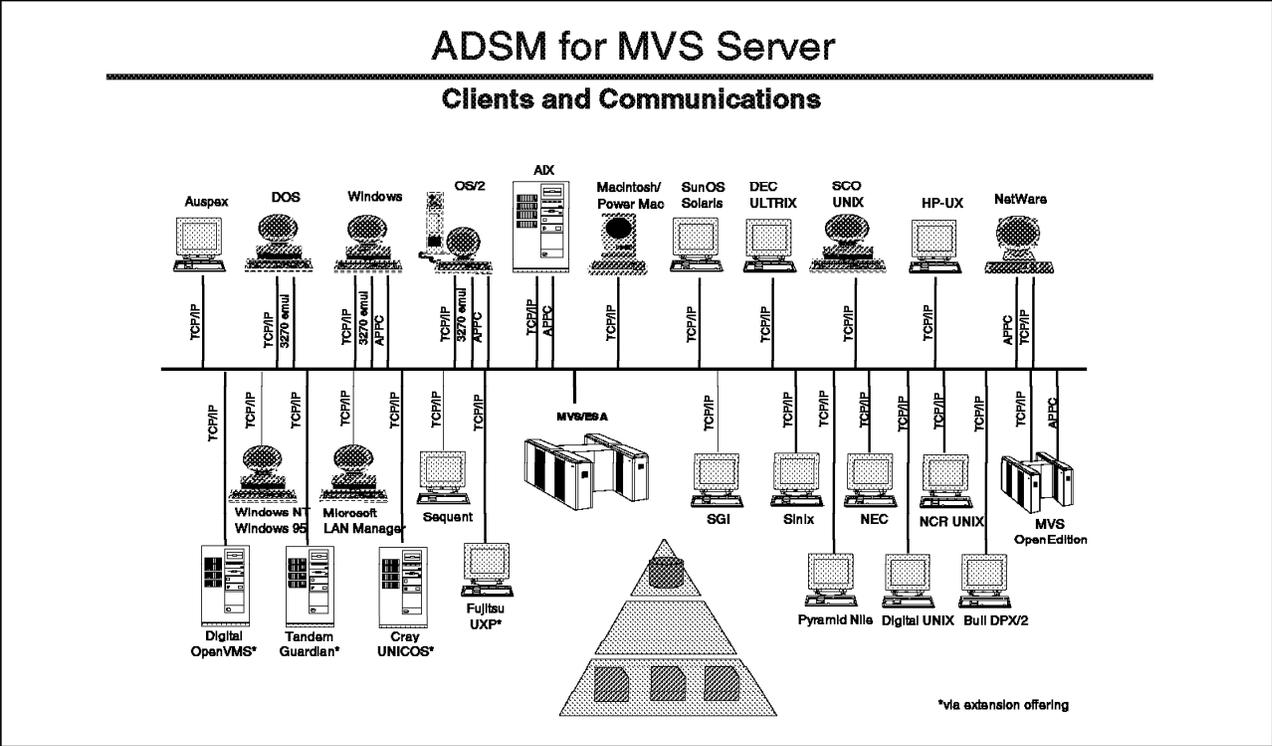
ADSM for MVS Server: Clients and Communications

This graphic shows the backup-archive clients and communication protocols that the ADSM for MVS server supports.

The flyers entitled *ADSM MVS/VM Fact Sheet* (G229-1267) and *ADSM for MVS Spec Sheet* (G225-6646) show the latest clients and communication protocols supported by ADSM for MVS. You can also get the most up-to-date information from the ADSM home page on the Web.

ADSM for MVS Server

Device Support

such as, but not limited to:

Any Disk Drives Supported by MVS

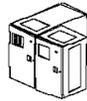1/2" Tape Drives
- IBM 3480
- IBM 3490/3490E
- IBM 3590 B11
- IBM compatible 3480, 3490, or 3490E tape drives

1/2" Tape Libraries
- IBM 3494
- IBM 3495
- Comparex 6388
- Memorex 5x00
- STK 4400, 9310, 9360
- EMASS AML

ADSM for MVS Server: Device Support

This graphic shows examples of devices that the ADSM for MVS server supports, including any disk storage device supported by MVS. It is not a complete list, and new products are continuously being added. For the most up-to-date list of supported devices, call the IBM Information Support Center in the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local IBM representative or reseller. You can also get the most up-to-date information from the ADSM home page on the Web. The flyer entitled *ADSM for MVS Supported Storage Devices* (G225-6619) is continuously updated and shows the latest supported devices.

The different UNIX clients require the UNIX Environment Support feature, and the desktop clients require the Desktop Environment Support feature. The OpenEdition MVS client requires the Open Systems Environment Support feature.

```
+------------------------------------------------------------------+
|                    ADSM for OS/2 Server                          |
| ================================================================ |
|                         Overview                                 |
|                                                                  |
|   ★ Operating System              ★ Functions Supported          |
|   – OS/2 Warp 3.0                 – HSM                           |
|   – Warp Connect                  – POR                          |
|   – Warp Server                                                  |
|   – Warp Server SMP               ★ Product Number               |
|                                   – 5801-AAR  (USA and Canada)   |
|   ★ Processor Support             – 5807-AAR  (additional for    |
|   – Any Level 386 or above Processor      Network Edition)       |
|                                   – 5639-A08  (others)           |
|   ★ Communication Protocols                                      |
|   – TCP/IP                                                        |
|   – Netware IPX/SPX                                               |
|   – APPC                                                          |
|   – NetBIOS                                                       |
|   – Named Pipes                                                   |
|   – Shared Memory                                                |
+------------------------------------------------------------------+
```

ADSM for OS/2 Server: Overview

The ADSM for OS/2 server is an ADSM Version 2.1 server.

Detailed information about the use of ADSM with OS/2 can be found in the redbooks entitled *ADSM for OS/2: Advanced Topics* (SG24-4740) and *Using ADSM to Back Up OS/2 LAN Server and Warp Server* (SG24-4682).

For the latest information about the ADSM for OS/2 server requirements or clients supported, see the ADSM Web page at:

- http://www.storage.ibm.com/storage/software/adsm/ad2serv.htm

Operating System

The ADSM for OS/2 Version 2.1 server runs on one of the following operating systems:

- OS/2 Warp 3.0, plus REXX

- Warp Connect with REXX

- Warp Server

- Warp Server SMP

Processor Support

The ADSM for OS/2 server runs on any Intel 386 or above processor, with at least 16 MB of memory.

Communication Protocols

The ADSM for OS/2 server supports the following communication protocols:

- TCP/IP for OS/2 3.0 or higher; comes with Warp Connect and Warp Servers

- IPX/SPX, available through NetWare Requester for OS/2 2.11, for Warp Connect and Warp Servers

- APPC (LU6.2), which comes with either IBM Communications Manager/2 1.0, or IBM SAA Networking Services/2

- NetBIOS, available through one of the following:

  – Warp Connect and Warp Servers

  – IBM Network Transport Services/2 (NTS/2)

  – IBM OS/2 LAN Server 2.0, 3.0, 4.0, Advanced Version 3.0 or 4.0

- Named Pipes, which comes with OS/2

- Shared Memory, which comes with OS/2

  Shared Memory is the only communication protocol supported in the stand-alone version.

Functions Supported

The ADSM for OS/2 server supports HSM for these client platforms:

- IBM AIX

- Sun Solaris 2.5

- Novell NetWare

- Windows NT

HSM clients for AIX and Solaris require the Space Management feature to be installed. The AvailHSM (formerly NetSpace) product is required on NetWare, and the OPEN/stor product is required on Windows NT. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

The ADSM for OS/2 server supports POR.

Product Numbers

In the USA and Canada, ADSM for OS/2 server is available as product number 5801-AAR for the Single Server Edition. For multiple users, the Network Edition Enabler, product number 5807-AAR, must also be ordered. For further ordering information, see the current announcement letter. Note that although AIX, OS/2, and Windows NT have the same order number, they have different feature numbers.

Outside the USA and Canada, the order number for the OS/2 server is 5639-A08.
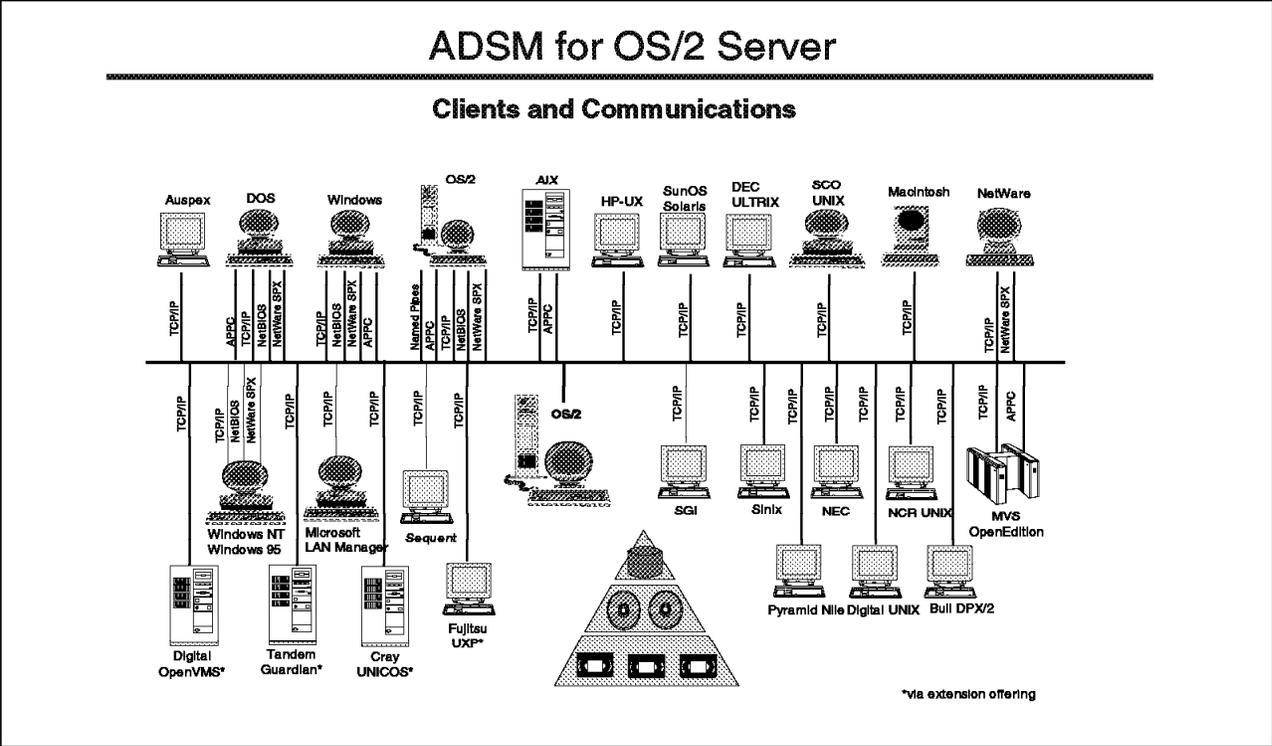
ADSM for OS/2 Server

Clients and Communications

ADSM for OS/2 Server: Clients and Communications

This graphic shows the backup-archive clients and communication protocols that the ADSM for OS/2 server supports.

The flyer entitled *ADSM/2 Spec Sheet* (G229-1269) shows the latest clients and communication protocols supported by ADSM for OS/2. You can also get the most up-to-date information from the ADSM home page on the Web.

ADSM for OS/2 Server

Device Support

such as, but not limited to:

Any Disk Drive Supported by OS/2

4 mm Tape Drives
– HP Jetstore
– IBM 3440
– Sony SDT
– WangDat

8 mm Tape Drives
– Andataco Encore
– Exabyte EXB
– IBM 3445, 3523
– SUN 8505
– Transitional Technologies CTS

QIC Tape Drives
– IBM 3450
– Tecmar QT525ES

DLT Drives
– Quantum DLT
– DEC DLT

Optical Drives
– IBM 3431
– ImageStar

4 mm Tape Libraries
– ADIC 1200D
– AIWA AL-D210
– HP SureStore
– Exabyte EXB
– StorageTek

8 mm Tape Libraries
– Andataco Encore
– Boxhill
– Exabyte EXB
– Lago LS/380L DataWheel
– StorageTek

DLT Libraries
– Quantum DLT
– DEC DLT

Optical Libraries
– IBM 3995

ADSM for OS/2 Server: Device Support

This graphic shows examples of devices that the ADSM for OS/2 server supports, including any disk storage device supported by OS/2. It is not a complete list, and new products are continuously being added. For the most up-to-date list of supported devices, call the IBM Information Support Center in the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local IBM representative or reseller. You can also get the most up-to-date information from the ADSM home page on the Web. The flyer entitled *ADSM/2 Device Support List* (G225-6621) is continuously updated and shows the latest supported devices.

The different UNIX clients require the UNIX Environment Support feature, and the desktop clients require the Desktop Environment Support feature. The Desktop environment is not required for the OS/2 client on this server. The OpenEdition MVS client requires the Open Systems Environment Support feature.

```
┌─────────────────────────────────────────────────────────────────────┐
│                    ADSM for Sun Solaris Server                        │
│  ─────────────────────────────────────────────────────────────────   │
│                            Overview                                   │
│                                                                       │
│                   ★ Operating System                                  │
│                   – Sun Solaris 2.5.1                                 │
│                                                                       │
│                   ★ Processor Support                                 │
│                   – Any processor with Sun4m or Sun4u support         │
│                                                                       │
│                   ★ Communication Protocol                            │
│                   – TCP/IP                                            │
│                                                                       │
│                   ★ Functions Supported                               │
│                   – HSM                                               │
│                   – DRM                                               │
│                   – POR                                               │
│                                                                       │
│                   ★ Product Number                                    │
│                   – 5801-AAR (US, AP, Canada)                         │
│                   – 5639-B22 (others)                                 │
│                                                                       │
└─────────────────────────────────────────────────────────────────────┘
```

ADSM for Sun Solaris Server: Overview

The ADSM for Sun Solaris server is an ADSM Version 2.1 server.

For the latest information about the ADSM for Sun Solaris server requirements or clients supported, see the ADSM Web page at:

- http://www.storage.ibm.com/storage/software/adsm/adsserv.htm

Operating System

The ADSM for Sun Solaris server runs on a Sun Solaris 2.5.1 system and must be installed on a file system. There is no raw logical volume support because there is no such concept with the base Solaris product.

Processor Support

To run the ADSM for Sun Solaris server, you must use a processor that supports the Sun4m or Sun4u architecture.

Communication Protocol

The ADSM for Sun Solaris server supports TCP/IP, which comes standard with Solaris.

Functions Supported

Functions Supported

The ADSM for Solaris server supports HSM for these client platforms:

- IBM AIX

- Sun Solaris 2.5

- Novell NetWare

- Windows NT

HSM clients for AIX and Solaris require the Space Management feature to be installed. The AvailHSM (formerly NetSpace) product is required on NetWare, and the OPEN/stor product is required on Windows NT. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

The ADSM for Solaris server supports POR.

Product Number
In the USA, Canada and Asia Pacific, the ADSM for Solaris server is available as product number 5801-AAR for one user. For multiple users, the user authorizations product number 5639-B22 must also be ordered. For further ordering information, look up the current announcement letter. Note that although AIX, OS/2, and Windows NT have the same order number, they have different feature numbers.

Outside the USA and Canada, the order number for the HP-UX server is 5639-B22.
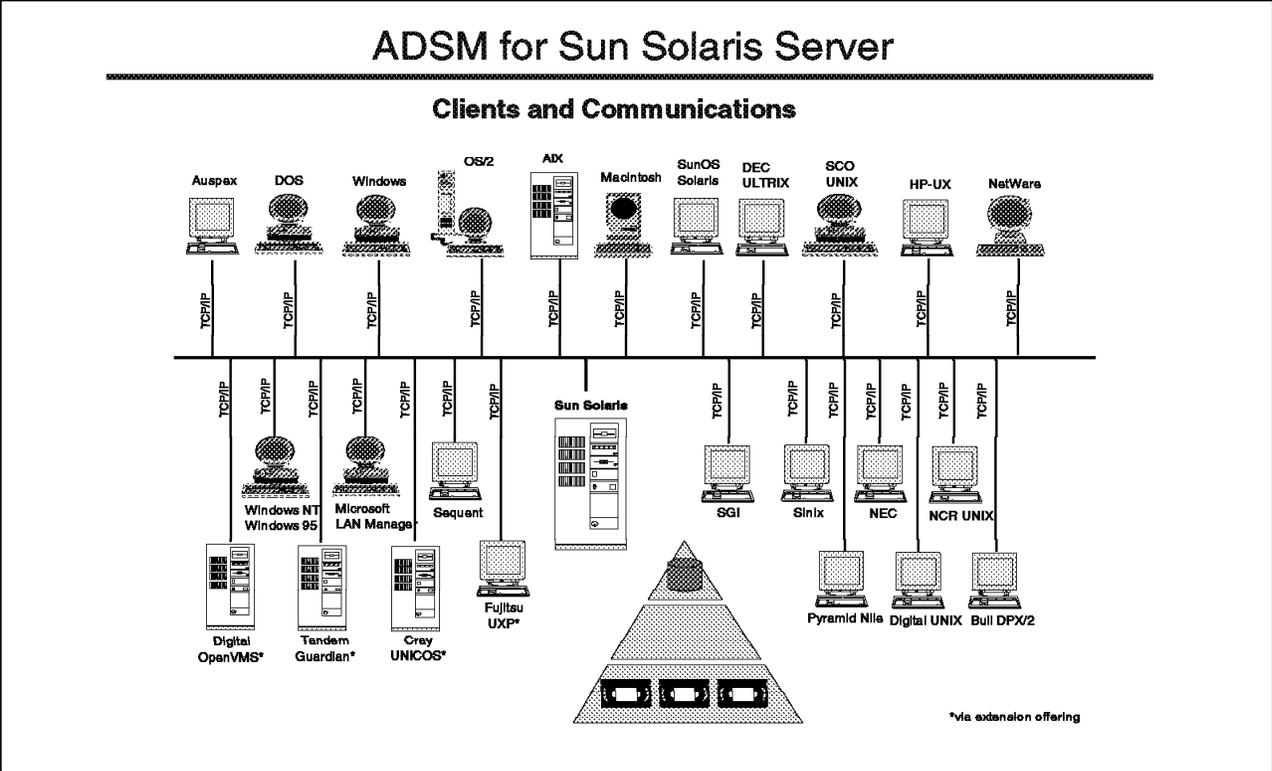
ADSM for Sun Solaris Server: Clients and Communications

This graphic shows the backup-archive clients and communication protocol that the ADSM for Sun Solaris server supports.

The flyer entitled *ADSM for Sun Solaris Fact Sheet* (G229-1272) shows the latest clients supported by ADSM for Sun Solaris. You can also get the most up-to-date information from the ADSM home page on the Web.

The different UNIX clients require the UNIX Environment Support feature, and the desktop clients require the Desktop Environment Support feature.

ADSM for Sun Solaris Server

Device Support

such as, but not limited to:

Any Disk Drives Supported by Sun Solaris 2.3

4 mm Tape Drives
- IBM 7206 Model 005
- HP C1520e, C1521e, C1533e
- HP Jetstore 2000e, 5000e, 6000e
- WangDat, DDS-2
- Sony SDT-5000, DDS-2

8 mm Tape Drives
- Andataco Encore 8205, 8505
- Dynatek HSB-2300, HSB-10.0, HSB-5000
- Exabyte EXB
- Transitional Technologies, Inc.
- SUN 8505XL

DLT Drives
- Quantum DLT 2000, 4000

1/2" Tape Drives
- IBM 3490E
- IBM 3590 Model B11

Optical Drives
- IBM 0632, IBM 7209

4 mm Tape Libraries
- Diverse Logistics Libra 8, Libra 16
- HP C1553A, C1561A
- HP SureStore 12000e
- IBM 7332-005
- ADIC VLS

8 mm Tape Libraries
- Andataco Encore
- Boxhill
- Exabyte EXB
- Lago LS/380L DataWheel
- Transitional Technologies, Inc. S10
- IBM 7331-205

1/2" Tape Libraries
- IBM 3494 L10, L12

DLT Libraries
- Quantum DLT 2500, 4500
- ATL Products

Optical Libraries
- IBM 3995
- HP 20T, 40T

ADSM for Sun Solaris Server: Device Support

This graphic shows examples of devices that the ADSM for Sun Solaris server supports, including any disk storage device supported by Sun Solaris 2.3. It is not a complete list, and new products are continuously being added. For the most up-to-date list of supported devices, call the IBM Information Support Center in the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local IBM representative or reseller. You can also get the most up-to-date information from the ADSM home page on the Web.

The flyer entitled *ADSM for Sun Solaris Device List* (G225-6626) is continuously updated and shows the latest supported devices.

```
                    ADSM for VM Server
━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━━
                         Overview


  ★ Operating System            ★ Functions Supported
  – VM/ESA V2.1.0               – HSM
  – VM/ESA V1.2.2               – POR
  – VM/ESA 1.1.5 370 Feature
                                ★ Product Number
  ★ Processor Support           – 5654-A02
  – Any S/390 Architecture Processor


  ★ Communication Protocols
  – TCP/IP
  – APPC
  – 3270
```

ADSM for VM Server: Overview

The ADSM for VM server is an ADSM Version 2.1 server.

The ADSM for VM server allows backups of the ADSM databases and storage pools to be directed to CMS flat files on VM minidisks or Shared File System directories when the device class is FILE. This sequential media can be used instead of tape for most ADSM functions, offering the advantage of avoiding intervention for tape mounts.

For the latest information about the ADSM for VM server requirements or clients supported, see the ADSM Web page at:

- http://www.storage.ibm.com/storage/software/adsm/ad3serv.htm

Operating System
The ADSM for VM Version 2.1 server runs on one of the following operating systems:

- VM/ESA Version 2 Release 1.0 or later

- VM/ESA Version 1 Release 2.2

- VM/ESA 1.1.5 370 feature

The following service levels are prerequisite:

- VMSES/E RSU9403 or above for VM/ESA Version 1

- VMSES/E RSU9601 or above for VM/ESA Version 2

- Server Virtual Machine (SVM) level 14

One of the following language run-time environments is also a prerequisite:

- IBM SAA AD/Cycle Language Environment/370 1.1 or later

- IBM C/370 Library Version 2 or later

- IBM Language Environment for MVS and VM Version 1.4 or later

Processor Support

ADSM for VM server runs on any S/390 architecture processor.

Communication Protocols

The ADSM for VM server supports the following protocols:

- TCP/IP, requiring either IBM TCP/IP Version 2 Release 3 for VM or later or IBM TCP/IP Version 2 Release 2 for VM with APAR PN35887

- APPC, requiring IBM ACF/VTAM 3.3 or later

- 3270, requiring IBM ACF/VTAM 3.3 or later

  ACF/VTAM is not required for the 3270 direct (non-SNA) attached protocol.

Functions Supported

The ADSM for VM server supports HSM for the following client platforms:

- IBM AIX

- Sun Solaris 2.5

- Novell NetWare

- Windows NT

HSM clients for AIX and Solaris require the Space Management feature to be installed. The AvailHSM (formerly NetSpace) product is required on NetWare, and the OPEN/stor product is required on Windows NT. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

The ADSM for VM server supports POR.

Product Number

The ADSM for VM Version 2.1 server is available as product number 5654-A02.
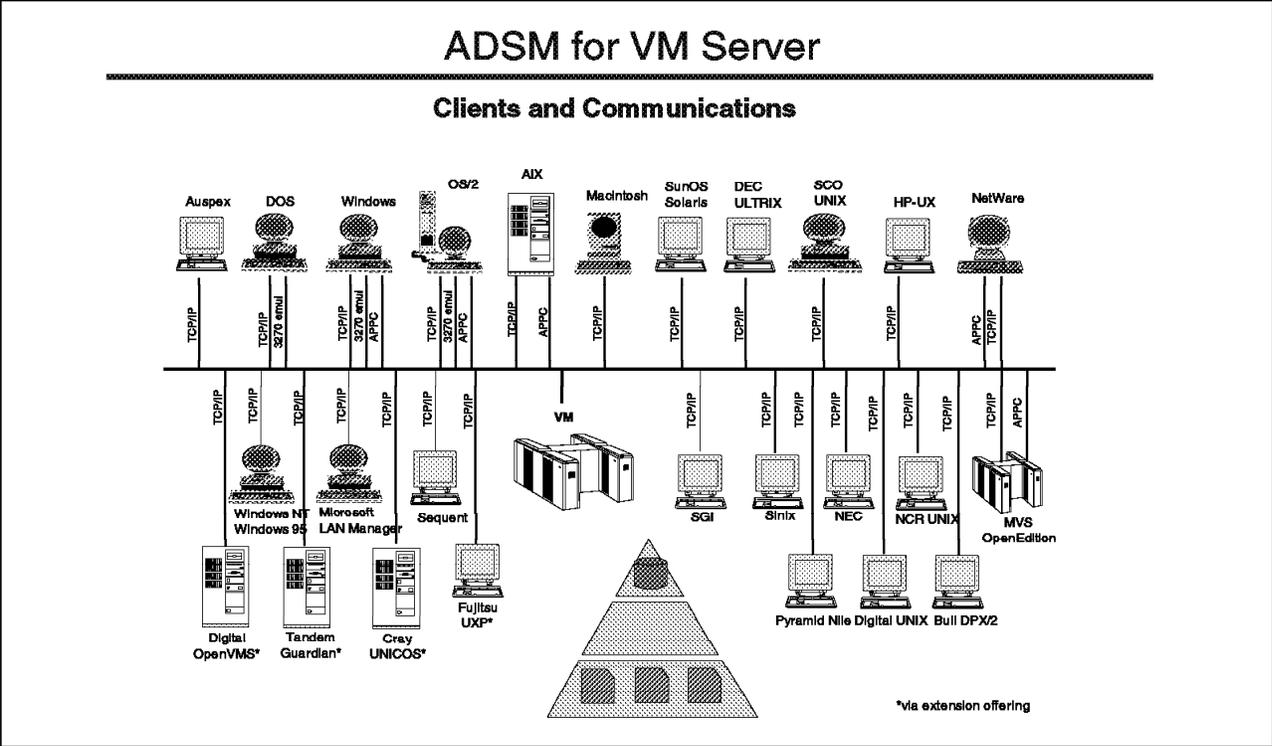
ADSM for VM Server

Clients and Communications

ADSM for VM Server: Clients and Communications

This graphic shows the backup-archive clients and communication protocols that the ADSM for VM server supports.

The flyers entitled *ADSM MVS/VM Fact Sheet* (G229-1267) and *ADSM for VM Spec Sheet* (G225-6645) show the latest clients and communication protocols supported by ADSM for VM. You can also get the most up-to-date information from the ADSM home page on the Web.

ADSM for VM Server

Device Support
such as, but not limited to:

Any Disk Drives Supported by VM

1/2" Tape Drives
– IBM 3420, 3422, 3430
– IBM 3480
– IBM 3490/3490E

1/2" Tape Libraries
– IBM 3494
– IBM 3495

ADSM for VM Server: Device Support

This graphic shows examples of devices that the ADSM for VM server supports, including any disk storage device supported by VM. It is not a complete list, and new products are continuously being added. For the most up-to-date list of supported devices, call the IBM Information Support Center in the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local IBM representative or reseller. You can also get the most up-to-date information from the ADSM home page on the Web. The flyer entitled *ADSM for VM Device Support List* (G225-6624) is continuously updated and shows the latest supported devices.

The different UNIX clients require the UNIX Environment Support feature, and the desktop clients require the Desktop Environment Support feature. The OpenEdition MVS client requires the Open Systems Environment Support feature.

```
                    ADSM for VSE/ESA Server
═══════════════════════════════════════════════════════════════════
                           Overview
                    ★ Operating System
                    – VSE/ESA Version 1.3 or later

                    ★ Processor Support
                    – Any VSE/ESA system

                    ★ Communication Protocols
                    – 3270 emulation
                    – APPC

                    ★ Functions Supported
                    – HSM

                    ★ Product Number
                    – 5686-073
```

ADSM for VSE/ESA Server: Overview

The ADSM for VSE/ESA server is an ADSM Version 1.2 server.

For details about the implementation of ADSM with VSE/ESA, see the redbook entitled *ADSM/VSE Implementation Guide* (SG24-4266).

For the latest information about the ADSM for VSE/ESA server requirements or clients supported, see the ADSM Web page at:

- http://www.storage.ibm.com/storage/software/adsm/advserv.htm

Operating System
The ADSM for VSE/ESA server runs on VSE/ESA Version 1.3 and requires the IBM C/370 Library 2.

Processor Support
The ADSM for VSE/ESA server runs on any processor supporting the VSE/ESA operating system.

Communication Protocols
The ADSM for VSE/ESA server supports 3270 emulation and APPC. Please note that there is no TCP/IP support in VSE/ESA and therefore no support for TCP/IP in the ADSM for VSE/ESA server. The clients connected to the ADSM for VSE/ESA server must be able to use 3270 emulation or APPC.

The 3270 emulation is included in VSE/ESA VTAM 3.4, which is a prerequisite for the VSE/ESA server.

Functions Supported

The ADSM for VSE/ESA server supports HSM for the Novell NetWare and Windows NT client platforms. AvailHSM (formerly NetSpace) on NetWare and OPEN/stor on Windows are required for this support. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

Product Number

The ADSM for VSE/ESA server is available as product number 5686-073.
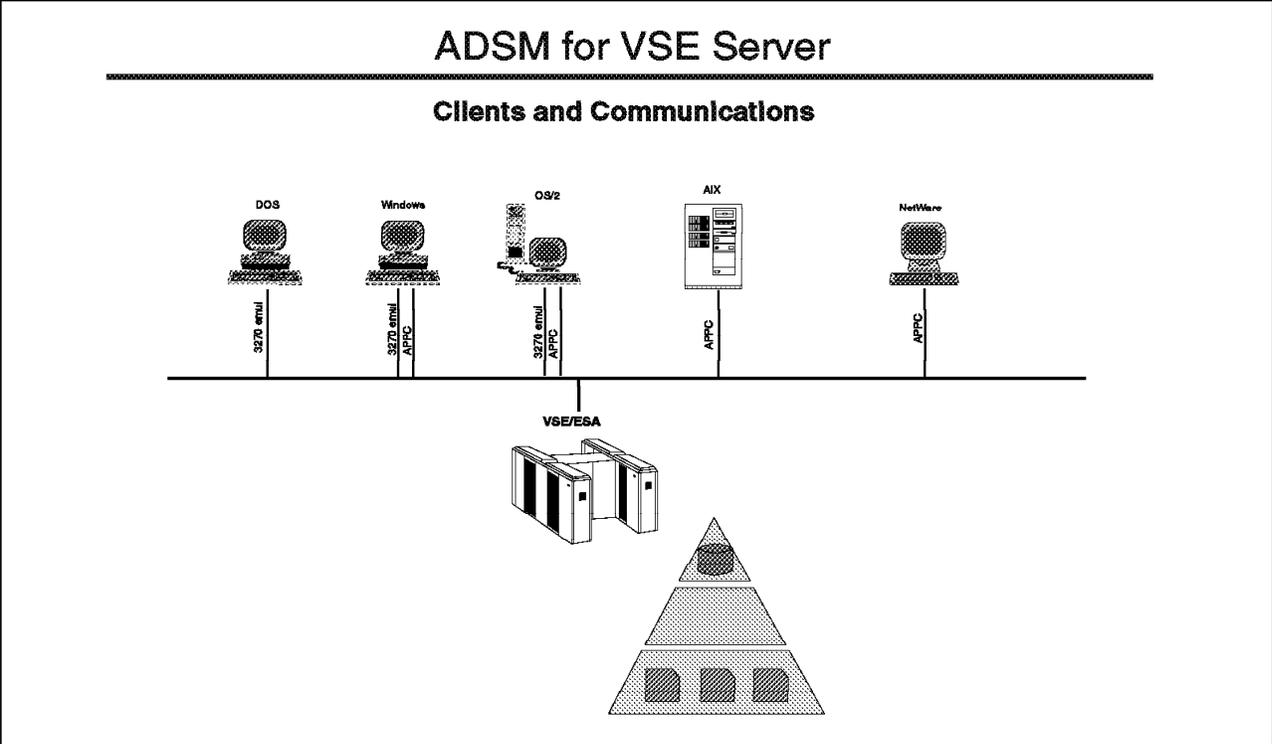
ADSM VSE/ESA Server: Clients and Communications

This graphic shows the backup-archive clients and communication protocols that the ADSM for VSE/ESA server supports. Please note that the ADSM for VSE/ESA server supports fewer backup-archive clients than the other ADSM server platforms because VSE/ESA does not support TCP/IP communications.

The flyer entitled *ADSM/VSE Fact Sheet* (G229-1271) shows the latest clients and communication protocols supported by ADSM for VSE/ESA. You can also get the most up-to-date information from the ADSM home page on the Web.

ADSM for VSE/ESA Server

Device Support

such as, but not limited to:

All DASD Supported by VSE/ESA
- RAMAC
- 3380
- 3390
- 9340

Tape Drives
- IBM Models 3480, 3490, 3490E

ADSM for VSE/ESA Server: Device Support

This graphic shows examples of devices that the ADSM for VSE/ESA server
supports, including any disk storage supported by VSE/ESA. For tape support,
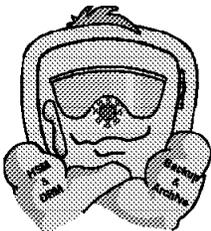3480, 3490, and 3490E devices are supported for parallel- or ESCON-channel
attachments. Optical devices and the IBM 3590 library are not supported.

This is not a complete list, and new products may be added. For the most
up-to-date list of supported devices, call the IBM Information Support Center in
the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local
IBM representative or reseller. You can also get the most up-to-date information
from the ADSM home page on the Web. The flyer entitled *ADSM for VSE/ESA
Device Validation List* (G225-6623) also shows the latest supported devices.

```
┌─────────────────────────────────────────────────────────────────────┐
│                                                                       │
│                  ADSM for Windows NT Server                           │
│         ──────────────────────────────────────────────────           │
│                         Overview                                      │
│                                                                       │
│               ★ Operating System                                      │
│               – Windows NT Workstation 3.51 or 4.0                    │
│               – Windows NT Server 3.51 or 4.0                         │
│                                                                       │
│               ★ Processor Support                                     │
│               – Any Intel 486 DX or Pentium Processor                 │
│                                                                       │
│               ★ Communication Protocols                               │
│               – TCP/IP                                                 │
│               – NetBIOS                                                │
│               – IPX/SPX                                                │
│               – Named Pipes                                           │
│                                                                       │
│               ★ Functions Supported                                   │
│               – HSM                                                    │
│               – DRM                                                    │
│                                                                       │
│               ★ Product Number                                        │
│               – 5801-AAR  (USA and Canada)                            │
│               – 5807-AAR  (additional for Network Edition)            │
│               – 5639-A09  (others)                                    │
│                                                                       │
└─────────────────────────────────────────────────────────────────────┘
```

ADSM for Windows NT Server: Overview

The ADSM for Windows NT server is an ADSM Version 2.1 server.

For detailed information on the ADSM for Windows NT server see a new redbook
called *ADSM Server for Windows NT Configuration and Recovery Examples*
(SG24-4878).

For the latest information about the ADSM for Windows NT server requirements
or clients supported, see the ADSM Web page at:

- http://www.storage.ibm.com/storage/software/adsm/adwserv.htm

Operating System
The ADSM for Windows NT server runs on either Windows NT Workstation 3.51
or 4.0 or Windows NT Server 3.51 or 4.0.

Processor Support
ADSM for Windows NT server runs on any Intel 486 DX or Pentium processor
with at least 32 MB of memory.

Communication Protocols
The ADSM for Windows NT server supports the following protocols, all of which
are included with the operating system:

- TCP/IP

- NetBIOS

- IPX/SPX

- Named Pipes

Functions Supported

The ADSM for Windows NT server supports HSM for these client platforms:

- IBM AIX

- Sun Solaris 2.5

- Novell NetWare

- Windows NT

HSM clients for AIX and Solaris require the Space Management feature to be installed. The AvailHSM (formerly NetSpace) product is required on NetWare, and the OPEN/stor product is required on Windows NT. Both products are supplied by Wang Software. For more information about AvailHSM, see "AvailHSM for Novell NetWare" on page 182, and for more information about OPEN/stor, see "OPEN/stor for Windows NT" on page 184.

The ADSM for Windows NT server supports the Disaster Recovery Manager feature, which is described in detail in Chapter 2, "Disaster Recovery Manager" on page 73.

Product Number

In the USA and Canada, ADSM for Windows NT Version 2.1 server is available as product number 5801-AAR for the Single Server Edition. For the Network Edition, the Network Edition Enabler, product number 5807-AAR, must also be ordered. For further ordering information, see the current announcement letter. Note that although AIX, OS/2, and Windows NT have the same order number, they have different feature numbers.

Outside the USA and Canada, the order number for the Windows NT server is 5639-A09.
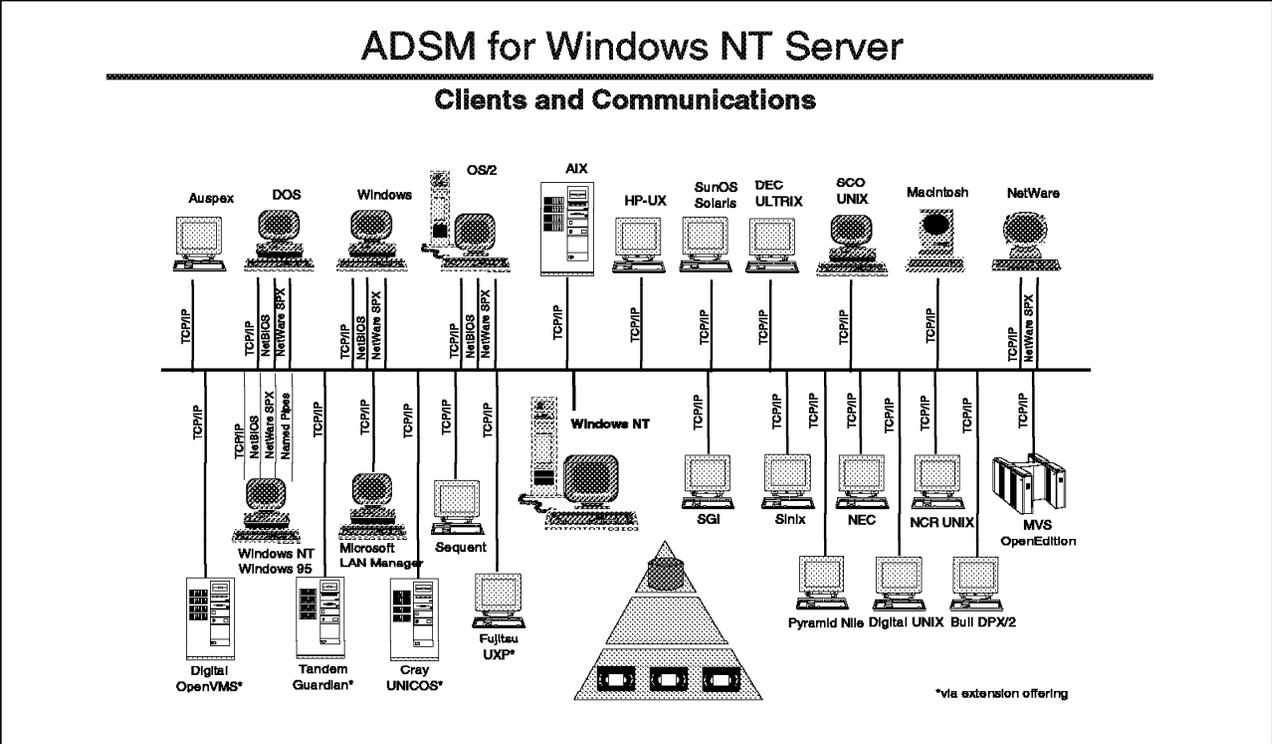
ADSM for Windows NT: Clients and Communications

This graphic shows the backup-archive clients and communication protocols that the ADSM for Windows NT server supports.

The flyer entitled *ADSM for Windows NT Version 2.1 Spec Sheet* (G225-6671) shows the latest clients and communication protocols supported by ADSM for Windows NT. You can also get the most up-to-date information on the Windows NT server from the ADSM home page on the Web.

The different UNIX clients require the UNIX Environment Support feature, and the desktop clients require the Desktop Environment Support feature. The Desktop environment is not required for the Windows NT client on this server. The OpenEdition MVS client requires the Open Systems Environment Support feature.

ADSM for Windows NT Server

Device Support

such as, but not limited to:

Any Disk Storage Supported by
Windows NT

4 mm Tape Drives
– HP Jetstore
– IBM 7206
– Sony SDT
– WangDat

8 mm Tape Drives
– Andataco Encore
– Exabyte EXB
– IBM 7208
– SUN 8505
– Transitional Technologies CTS

QIC Tape Drives
– IBM 7207

DLT Drives
– Quantum DLT
– DEC DLT

4 mm Tape Libraries
– ADIC 1200D
– AIWA AL-D210
– HP SureStore
– Exabyte EXB
– StorageTek
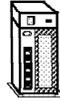
8 mm Tape Libraries
– Andataco Encore
– Boxhill
– Exabyte EXB
– Lago LS/380L DataWheel
– IBM 7331
– StorageTek

DLT Libraries
– Quantum DLT
– DEC DLT
– Odetics ACL

ADSM for Windows NT: Device Support

This graphic shows examples of devices that the ADSM for Windows NT server
supports, including any disk storage device supported by Windows NT. It is not
a complete list, and new products are continuously being added. For the most
up-to-date list of supported devices, call the IBM Information Support Center in
the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local
IBM representative or reseller. You can also get the most up-to-date information
from the ADSM home page on the Web. The flyer entitled *ADSM for Windows NT
Supported Storage Devices* (G225-6672) is continuously updated and shows the
latest supported devices.

Appendix B. Client Details

ADSM Clients: Overview

```
                    AIX                    OpenEdition MVS
                    Apple Macintosh        OS/2
                    Auspex                 Pyramid Nile
                    Bull DPX/2             SCO UNIX 386
                    CMS                    Sequent PTX
                    Digital ULTRIX         Sinix RISC
                    Digital UNIX           Sinix 386/486
                    DOS                    SGI IRIX
                    HP-UX                  SUN Solaris
                    NCR UNIX SVR4          Windows
                    NEC EWS-UX/V           Windows 32 bit
                    Novell NetWare

                    Special Support:

                    AS/400 API             Fujitsu UXP
                    CRAY UNICOS            Lotus Notes Backup Agent
                    Digital OpenVMS        Tandem Guardian
                                           WebShell
```

ADSM Clients: Overview

Different types of client and server platforms and different communication protocols can be used in the ADSM environment. Each platform supports a wide variety of storage devices and connectivity options.

ADSM supports a range of IBM and non-IBM client platforms for backup and archive. These include UNIX platforms such as AIX, Sun Solaris, and HP-UX, desktop platforms such as OS/2, DOS, and Windows, and the OpenEdition MVS client. Details about the clients can be found on the ADSM home page on the Web, or directly on the ADSM client platforms page:

• http://www.storage.ibm.com/storage/software/adsm/adclient.htm

Basically all clients are supported by all server platforms, independent of ADSM server Version 1 or 2. A number of communication protocols are supported for connectivity to the ADSM servers. What differs is which communication protocols are supported between a certain client and a certain server. The protocols supported depend on each client platform. This may limit the clients supported for a certain server. Refer to "ADSM Server Family" on page 245 for the server support.

The pages that follow give an overview of each client available and describe the hardware, operating system, and communications protocols they support. The clients are listed alphabetically; a section at the end deals with "special support" clients.

Most clients are provided with both a GUI and a CLI. The only exceptions are Macintosh, which has only a GUI, and Novell NetWare, which has only a CLI.

Cross-client restore and cross-file system restore are supported between UNIX clients or between Intel PC clients, but with some restrictions.

Not all UNIX systems implement ACLs in the same way, so this information is not preserved with cross-platform restores in UNIX environments. Both the format and syntax of extended ACL data may vary and thus could be lost when restored to a different operating system. Also be aware that even when you restore from, say, AIX to AIX, the ACL data will be restored with the original userids. So if two different users on the two systems have the same userid, permissions could be granted to an unauthorized user. If you use cross-client restore with ACL data, make sure the userids are properly defined.

ADSM provides the ability to back up files from a DOS, Windows, OS/2, or Windows 32 bit client and restore them to any of these clients. The files can also be restored to a different file system from which they were backed up. If you back up files with long file names, using an OS/2 or Windows 32 bit client, the file names are truncated when restored to a DOS or Windows client. Also, DOS and Windows clients cannot back up remote NT drives with volume labels longer than 12 characters. This is a DOS restriction.

If the files are backed up from an ADSM Version 2 client, they must be restored to a Version 2 client. Version 1 client data can be restored to either a Version 1 or Version 2 client. New attribute information is associated with each Version 2 client file during the backup, and only Version 2 clients recognize the information.

Please note that once you have started using a Version 2 client on a particular node, you will not be able to use a Version 1 client on that node. Data backed up with Version 2 clients cannot be restored by Version 1 clients, regardless of the server version.

The PTF README files for the clients list new features and provide advice for the installation. The README file will detail what is specific for that client both functionally and in the installation process. You should always use this README information.

AIX Client

- **Operating System**
  - ✓ **IBM AIX 3.2.3 and Later**
  - ✓ **IBM AIX 4.1 and Later**
- **Communication Protocols**
  - ✓ **TCP/IP (Standard)**
  - ✓ **APPC (through AIX SNA Server/6000 V2.1 or 3.1)**
  - ✓ **Shared Memory**
- **Hardware Supported**
  - ✓ **RS/6000 Workstations**
  - ✓ **Bull DPX/20 Workstations**
- **File System Considerations**

AIX Client

AIX supports backup-archive client, administrative client, and Space Management client CLIs and GUIs. Applications using the ADSM API are also supported. Space Management and administrative GUI clients require the AIX Windows Environment to be installed. There are separate client PTFs for AIX V3.2, AIX V4.1, and AIX V4.2.

For the use of AIX clients with ADSM, refer to the redbook entitled *Getting Started with ADSM AIX Clients* (GG24-4243).

Operating System
The AIX client is supported by AIX Version 3.2.3 or later and AIX Version 4.1 or later.

Communication Protocols
The AIX client can use one of the following communication protocols: either TCP/IP, which comes standard with AIX/6000, or APPC (LU6.2), which is available through AIX SNA Server/6000 Version 2.1 for AIX Version 3.2, or AIX SNA Server/6000 Version 3.1 for AIX Version 4.1, or the Shared Memory protocol, if the ADSM server and client are installed on the same machine.

Hardware Supported

Any RS/6000 with at least 16 MB memory

The AIX client also provides support for Bull AIX environment on the Bull DPX/20 systems, with the same functions as for the RS/6000 environment.
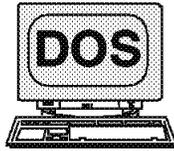
File System Considerations

ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- Network File System (NFS)
- Andrew File System (AFS)
- Distributed File System (DFS)

ADSM backs up regular files and directories. It does not back up character special files, block special files, or first-in-first-out (FIFO) special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files as well as the ACLs.

The AIX client supports backup of AFS and DFS files, ACLs, and mount points. See Chapter 4, "Andrew File System and Distributed File System Support" on page 103 for more details.

Apple

- **Operating System**
  - ✓ **Macintosh Operating System 6.0.7 or System 7**
- **Communication Protocols**
  - ✓ **MacTCP 2.0.2 through 2.0.6**
- **Hardware Supported**
  - ✓ **Apple Macintosh or Apple Macintosh Power PC**
- **File System Support**

Apple Macintosh Client

The Macintosh platform does not have an administrative client. To get an administrative client you must use one of the other backup-archive clients that have this support.

This is a GUI-only platform; there is no command line support.

Operating System
The Macintosh ADSM client requires the Macintosh Operating System 6.0.7 or System 7.

Communication Protocols
The Macintosh client uses TCP/IP, which is available through MacTCP Versions 2.0.2—2.0.6.

Hardware Supported
The Macintosh client runs on an Apple Macintosh or an Apple Macintosh PowerPC.

File System Support
The Macintosh file structure is based on volumes, folders, and files. It is really a hierarchical structure, where the folders are the directories. The files can be as big as the volume, which is up to 2GB.

When you back up files, ADSM also backs up related file and folder attributes, including AppleShare privileges.

Auspex Client

```
┌─────────────────────────────────────────────────────────────┐
│                      Auspex Client                          │
│  ═══════════════════════════════════════════════════════    │
│   ┌───────────────────────────────────────────────────┐     │
│   │  ⊛  A U S P E X   S Y S T E M S                    │     │
│   └───────────────────────────────────────────────────┘     │
│                                                             │
│           ▪ Operating System                                │
│              ✓ Solaris 2.3 or 2.4                           │
│           ▪ Communication Protocols                         │
│              ✓ TCP/IP (Standard)                           │
│           ▪ Hardware Supported                             │
│              ✓ Auspex NFS File Server Based on             │
│                Sun Processor                               │
│           ▪ Miscellaneous                                 │
│              ✓ Support Provided by SunOS 4.1.3 Client      │
│           ▪ File System Considerations                    │
│                                                             │
└─────────────────────────────────────────────────────────────┘
```

Auspex Client

Auspex NFS file servers are built on Sun processors and use the SunOS 4.1.3 operating system. The Auspex client supports backup-archive client and administrative client CLIs and GUIs. Applications using the ADSM API are also supported.

Operating System

The Auspex client uses SunOS 4.1.3. See also "Sun Solaris Client" on page 310 for details about the Sun Solaris client.

Communication Protocols

The Auspex client uses TCP/IP, which comes standard with SunOS.

Hardware Supported

Any Auspex NFS file server built on Sun processors.

File System Considerations

ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

Bull DPX/2 Client

Worldwide
Information
Systems
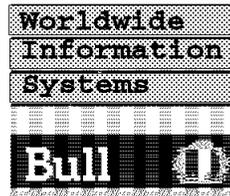
Bull

- **Operating System**
    - ✓ **DPX/2 300 B.O.S. Release 2.00**
- **Communication Protocol**
    - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
    - ✓ **Bull DPX/2 300 Workstation with**
        68000 Processor
- **File System Considerations**

BULL DPX/2 Client

Bull systems (for example, Bull DPX/20) with AIX operating systems can be ADSM clients if you install the ADSM AIX client support. The Bull DPX/2 client has backup-archive client and administrative client CLIs and GUIs. Applications using the ADSM API are also supported.

Operating System

The Bull DPX/2 client runs on Bull Open Software (BOS) Release 2.

Communication Protocols

The Bull DPX/2 client uses TCP/IP, which comes standard with the INET software package.

Hardware Supported

The Bull DPX/2 client runs on the Bull DPX/2 300 family workstation with a 68000 processor.
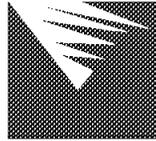
File System Considerations

ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

CMS Client
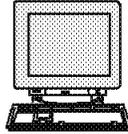
- **Operating System**
  - ✓ **VM/ESA Version 2.1.0 and Later**
  - ✓ **VM/ESA Version 1.2.2**
  - ✓ **VM/ESA 1.1.5 370 Feature**
- **Communication Protocols**
  - ✓ **TCP/IP**
    IBM TCP/IP Version 2.3 for VM
    IBM TCP/IP Version 2.2 (APAR PN35887 required)
  - ✓ **APPC**
  - ✓ **3270 Communication Protocol**
- **Hardware Supported**
  - ✓ **Any S/390 Processor**

CMS Client

There is no actual ADSM CMS backup-archive client. There is, however, an administrative client that is part of the VM server. This administrative client has the same systems requirements as the VM server. See "ADSM for VM Server: Overview" on page 271 for more details.

Operating System
The CMS administrative client uses VM/ESA Version 2.1.0 and later, Version 1.2.2, or VM/ESA 1.1.5 370 Feature. Certain service levels and specific language environments are required.

Communication Protocols
The CMS administrative client uses TCP/IP, APPC, or 3270 communication protocols. TCP/IP requires IBM TCP/IP Version 2.3 for VM or later, or IBM TCP/IP Version 2.2 with APAR PN35887. APPC and 3270 require ACF/VTAM Version 3.3 or later.

Hardware Supported
Any S/390 architecture processor

Digital Ultrix Client

digital

- **Operating System**
  - ✓ **Digital Ultrix 4.2A, 4.3, 4.4 or 4.5**
- **Communication Protocol**
  - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
  - ✓ **Any DECstation**
- **File System Considerations**

Digital Ultrix Client

The Digital Ultrix client was formerly known as DEC Ultrix. It provides a backup-archive CLI and GUI an administrative client CLI.

Operating System

The Digital Ultrix client supports Digital Ultrix 4.2A , 4.3, 4.4, and 4.5.

Communication Protocols

The Digital Ultrix client uses TCP/IP, which comes standard with Digital Ultrix.

Hardware Supported

The Digital Ultrix client runs on any DECstation.

File System Considerations

ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

Digital UNIX Client

digital

- **Operating System**
  - ✓ **Digital UNIX V3.2 (DEC OSF/1)**
- **Communication Protocol**
  - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
  - ✓ **DEC2000, 3000, 4000, 7000, or 10000**
  - ✓ **Alphaserver 1000, 2000, 2100, 2100A,**
       8200, or 8400
- **File System Considerations**

Digital UNIX Client

ADSM provides a client for Digital UNIX (formerly known as DEC OSF/1). It provides a backup-archive CLI and GUI an administrative client CLI. Applications using the ADSM API are also supported.

Operating System
The Digital UNIX client supports Digital UNIX Version 3.2.

Communication Protocols
The Digital UNIX client uses TCP/IP, which comes standard with Digital UNIX.

Hardware Supported
The Digital UNIX client runs on the Alphaserver 1000, 2000, 2100, 2100A, 8200, or 8400, or on the DEC2000, 3000, 4000, 7000, or 10000.

File System Considerations
ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

```
                            DOS Client
  ─────────────────────────────────────────────────────────────

                              ┌─────────┐
                              │  DOS    │
                              └─────────┘

                    ▪ Operating System
                       ✓ DOS 5.0 or Later
                    ▪ Communication Protocols
                       ✓ TCP/IP
                       ✓ IPX/SPX
                       ✓ NetBIOS
                       ✓ 3270
                       ✓ APPC
                    ▪ Hardware Supported
                       ✓ Any PC
```

DOS Client

The DOS client provides a backup-archive client CLI and GUI and an
administrative client CLI. DOS also supports double-byte character set (DBCS)
for Japanese and simplified Chinese.

Operating System

The DOS client requires DOS 5.0 or later. DBCS requires PC DOS J6.1/V for
Japanese, and PC DOS P6.3/V for simplified Chinese.

Communication Protocols

The DOS client requires any one of these communication protocols:

- IPX/SPX, available through NetWare Requester for DOS 3.01

- NetBIOS, available through IBM LAN Support Program

- TCP/IP, available through IBM TCP/IP 2.0 or 2.1 for DOS or FTP Software
  PC/TCP 3.0 for DOS

- 3270, available through IBM Personal Communications/3270 for DOS Version
  3.0 or Attachment EXTRA! for DOS Version 2.3

- APPC, available through IBM PC Support/400 2.3 or later, and IBM LAN
  Support Program 1.2.6 or later, or through IBM Client Access/400 for
  Extended DOS 3.0 or later

The DBCS support requires either 3270, which is available through IBM
Japanese 3270 PC, or TCP/IP, which is available through IBM TCP/IP 2.0 or 2.1
for DOS (AP kit).

Hardware Supported

The DOS client requires a personal computer with at least 512 KB memory.
DBCS requires an IBM PS/55 with 4 MB of memory.

HP-UX Client

- **Operating System**
  - ✓ **HP-UX 8.07, 9.01, 9.03, 9.04, or 10.0**
- **Communication Protocol**
  - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
  - ✓ **Any HP 9000 Series 700 or 800**
- **File System Considerations**

HP-UX Client

The HP-UX client provides backup-archive client and and administrative client CLIs and GUIs. Applications using the ADSM API are also supported. There are now separate PTFs for the HP-UX V8/9 and the HP-UX V10. The HP-UX V10 client also supports the X/Open API in PTF 6.

Operating System
The HP-UX client requires HP-UX level 8.07, 9.01, 9.03, 9.04, or 10.0. The administrative client GUI works with HP-UX level 9.03, 9.04, or 10.20 and requires an X Windows System 11.5 and Motif 1.2 Windows Manager.

Communication Protocol
The HP-UX client supports TCP/IP, which comes standard with HP-UX.

Hardware Supported
The HP-UX client supports any HP 9000 series 700 or 800 workstation.

File System Considerations
ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

NCR UNIX SVR4 Client

NCR

- **Operating System**
  - ✓ **NCR UNIX SVR4 MP-RAS Release 2.02,**
    2.03, or 3.0
- **Communication Protocol**
  - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
  - ✓ **NCR WorldMark System 3000 Series**
    (Including Model 5100)
- **File System Considerations**

NCR UNIX SVR4 Client

The NCR UNIX SVR4 (formerly known as AT&T GIS SVR4) client provides a backup-archive client CLI and GUI and an administrative client CLI. Applications using the ADSM API are also supported.

Operating System
The NCR UNIX client requires NCR UNIX SVR4 MP-RAS Release 2.02, 2.03, or 3.0.

Communication Protocol
The NCR UNIX client supports TCP/IP, which comes standard with NCR UNIX SVR4 MP-RAS.

Hardware Supported
The NCR UNIX client supports any NCR WorldMark System 3000 Series, including the Model 5100.

File System Considerations
ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

NEC EWS-UX/V Client

NEC

- **Operating System**
  - ✓ **NEC EWS-UX/V Release 4.2,**
    Revisions 9.1, 9.2, 10.1, or 11.1
- **Communication Protocol**
  - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
  - ✓ **NEC EWS 4800 Series Model 320SX**
- **File System Considerations**

NEC EWS-UX/V Client

Client support is provided for the Nippon Electronic Corporation (NEC) UNIX platform, EWS-UX/V. It provides a backup-archive CLI and GUI and an administrative client CLI. Applications using the ADSM API are also supported.

Operating System
The NEC EWS-UX/V client requires NEC EWS-UX/V Release 4.2, Revisions 9.1, 9.2, 10.1, or 11.1.

Communication Protocol
The NEC EWS-UX/V client uses TCP/IP, which comes standard with the NEC EWS-UX/V operating system.

Hardware Supported
An NEC EWS 4800 series Model 320SX workstation is required to run an ADSM NEC EWS-UX/V client.

File System Considerations
ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

Novell NetWare Client

NetWare

- **Operating System**
  - ✓ **Novell NetWare 3.11, 3.12, 4.01, 4.02, 4.1, or 4.1.1**
- **Communication Protocols**
  - ✓ **TCP/IP (Standard)**
  - ✓ **IPX/SPX (Standard)**
  - ✓ **APPC (through NetWare for SAA)**
- **Hardware Supported**
  - ✓ **Any PC Running a NetWare File Server**
- **File System Considerations**

Novell NetWare Client

The Novell NetWare platform supports only the backup-archive client CLI. There is no support for administrative clients. To get an administrative client you must use one of the other backup-archive clients that have this support. Applications using the ADSM API are also supported.

For the use of NetWare clients with ADSM, refer to the redbook entitled *Getting Started with ADSM NetWare Clients* (GG24-4242).

Operating System
The Novell NetWare client supports Novell NetWare servers at level 3.11, 3.12, 4.01, 4.02, or 4.1.

Communication Protocols
The Novell NetWare client uses any of the following protocols:

- TCP/IP, which comes standard with Novell NetWare
- IPX/SPX, which comes standard with Novell NetWare
- APPC (LU6.2), which is available through NetWare for SAA

For Version 3.11 only, PWSCS is supported through the IBM VM Programmable Workstation Communication Services 1.1 or later.

Hardware Supported

Any personal computer or workstation running as a NetWare file server

File System Considerations

The Novell NetWare server has its own file system based on server name, volume, directories, and file. It uses a bindery as an information database. The bindery contains definitions for users and groups and is three files grouped together as one database. ADSM backs up the bindery information as well as the files. In NetWare Version 4.01 the bindery concept has been replaced with the Netware Directory Service (NDS). ADSM backs up the bindery information and the NDS.

The NetWare file system maintains both FAT information and a separate Directory Entry Table (DET). With attached DOS, Macintosh, or OS/2 HPFS machines, NetWare provides a file name mapping facility stored in namespaces. ADSM backs up these namespaces.

Operating System

  ✓ **MVS/ESA SP V5R1 or Later with**

    SMP/E Release 8.0 or Later

Communication Protocols

  ✓ **TCP/IP Release 3.1 for MVS**

  ✓ **APPC (through ACF/VTAM Release 4.2 or Later)**

Hardware Supported

  ✓ **Any S/390 Processor**

File System Considerations

OpenEdition MVS Client

Backup-archive support is provided for the OpenEdition MVS UNIX file system, also known as the Hierarchical File System (HFS). It provides a backup-archive CLI and GUI and an administrative client CLI. Applications using the ADSM API are also supported. There is a separate PTF for the API client.

Operating System
The OpenEdition MVS client requires MVS/ESA SP V5R1 or later, with SMP/E Release 8.0 or later.

Communication Protocols
The OpenEdition MVS client requires either TCP/IP Release 3.1 for MVS or APPC (LU6.2), which is available through ACF/VTAM Release 4.2 or later.

The client GUI function is provided through X Windows, which requires TCP/IP.

Hardware Supported
Any System/390 processor with the proper operating system level is supported.

File System Considerations
The UNIX file system is a hierarchical file system of directories, subdirectories, and files. OpenEdition MVS is compliant with the POSIX 1003.1 standard and contains the same file system structure as other UNIX systems. Every file system is stored in an HFS data set, which can be up to 2 GB and contains a set of files and directories.

The HFS is a special MVS partitioned data set extended (PDSE) that is allocated on SMS-managed DASD. This special PDSE can be mounted as a file system

and can also be accessed by standard MVS batch programs. Please note that the ADSM OpenEdition MVS client provides file level backup-archive and is supported with ADSM V2 servers only. Because of differences in code pages, it is not possible to perform ADSM cross-platform restore. You can, however, restore across different ADSM OpenEdition MVS clients.

You can also back up an OpenEdition MVS HFS by NFS mounting it to an ADSM client that has NFS support. This is possible because the OpenEdition MVS NFS server feature supports the HFS. NFS client support is not yet available on OpenEdition MVS, so you cannot mount remote file systems to OpenEdition MVS and back them up there.

When mounting the HFS to an NFS client, you must consider that MVS is EBCDIC. You can either mount the HFS in binary (EBCDIC) or text (ASCII) modes. You may want to mount in binary mode if, for example, you are transferring tar or binary files from one OpenEdition MVS system to your workstation and then to another OpenEdition MVS system. OpenEdition MVS provides a command to convert ASCII files to EBCDIC. If you mount the HFS in text mode, NFS handles the conversion.

OS/2 Client

- **Operating System**
  - ✓ **IBM OS/2 2.0, 2.1, 2.11**
  - ✓ **OS/2 Warp 3.0 or Warp Server**
  - ✓ **IBM OS/2 LAN Server 3.0 or 4.0**
- **Communication Protocols**
  - ✓ **TCP/IP**
    - **IBM TCP/IP 2.0 for OS/2 2.0**
    - **IBM TCP/IP 3.0 for OS/2 Warp (or Warp Connect) 3.0**
  - ✓ **IPX/SPX**
    - **NetWare Requester for OS/2 2.01 or OS/2 2.11**
  - ✓ **NetBIOS**
    - **IBM NTS/2**
    - **IBM OS/2 LAN Server 2.0, 3.0, 4.0, Advanced Version 3.0 or 4.0**
- **Hardware Supported**
  - ✓ **Any PC That Supports OS/2**
- **File System Considerations**

OS/2 Client

The OS/2 client supports backup-archive and administrative client CLIs and GUIs. Applications using the ADSM API are also supported. OS/2 also supports DBCS for Japanese and simplified Chinese.

Refer to "Lotus Notes Backup Agent" on page 320 for more information.

For the use of OS/2 with the ADSM client, refer to the redbook entitled *Using ADSM to Back Up OS/2 LAN Server and Warp Server* (SG24-4682).

Operating System

The OS/2 client requires one of the following OS/2 environments:

- OS/2 Version 2.0, 2.1, or 2.11

- OS/2 Warp Version 3.0 or Warp Server

- OS/2 LAN Server 3.0 or 4.0 Advanced, with 386 HPFS installed

The administrative GUI or the API works only with OS/2 2.0, 2.1, and 2.11 or OS/2 Warp 3.0. DBCS requires OS/2 Warp J 3.0 for Japanese and OS/2 Warp P 3.0 for simplified Chinese.

Communication Protocols

The OS/2 client requires any one of these communication protocols:

- TCP/IP, available through IBM TCP/IP 2.0 for OS/2 2.0 or IBM TCP/IP 3.0 for OS/2 Warp (or Warp Connect) 3.0 or Warp Server

- IPX/SPX, available through NetWare Requester for OS/2 2.01 (for OS/2 2.0 or 2.1), NetWare Requester for OS/2 2.11 (for OS/2 Warp), or OS/2 Warp Connect or Warp Server

- NetBIOS, available through IBM OS/2 LAN Server 2.0, 3.0, or 4.0 Advanced Version 3.0 or 4.0, IBM Network Transport Services (NTS/2), or OS/2 Warp Connect or Warp Server

  Note that Novell NetBIOS emulation is not supported.

- 3270, available through IBM Communications Manager/2 1.0

- APPC, available through IBM Communications Manager/2 1.0, IBM SAA Networking Services/2, or IBM Client Access/400.

- Named Pipes, which comes with OS/2

- Shared Memory, which comes with OS/2

  Shared Memory is for Single Server only and is the only protocol that is supported on the stand-alone version. It is shipped with the client in the OS/2 2.1 Single Server Edition.

DBCS support requires either TCP/IP, which is available through IBM TCP/IP 2.0 or 2.1 with AP kit, or APPC, which is available through IBM Communications Manager/2 1.0 or IBM SSA Networking Services/2.

Hardware Supported

The OS/2 client runs on any personal computer supporting OS/2. DBCS requires an IBM PS/55 with 4 MB of memory.

File System Considerations

The OS/2 file system can be a FAT structure such as DOS, or the special OS/2 file systems HPFS and 386 HPFS. The HPFS structure is internally similar to FAT but with more attributes and higher function. ADSM backs up files and their attributes.

On OS/2 LAN Server, ADSM backs up local 386 HPFS files and directories and the access control profiles (ACPs) stored within them. To back up ACPs for FAT and HPFS files and directories, you must back up the NET.ACC or, preferably, the NETACC.BKP file. ADSM does not back up ACPs for FAT and HPFS files and directories at an individual file and directory level. You can also use the LAN Server BACKACC and RESACC utilities to back up and restore LAN Server ACPs.

On OS/2 Warp Server, ADSM supports FAT, HPFS, 386 HPFS, extended attributes, long file names, and ACPs.

OS/2 Warp Server provides local backup through OS/2 Warp Server backup/restore. For more information, see "OS/2 Warp Server Backup/Restore and ADSM" on page 180 and *Using ADSM to Back Up OS/2 LAN Server and Warp Server (SG24-4682).*

Please refer to the

Sinix RISC client

Pyramid Nile Client

The Pyramid Nile client is available through the ADSM Siemens Nixdorf Sinix RISC client. It supports a backup-archive client CLI and GUI and an administrative client CLI. See the "Sinix RISC Client" on page 307 for more information on Sinix.

Operating System

The Pyramid Nile client requires DC/OSx Level 79 or later, or Reliant UNIX 1.0 or later.

Communication Protocol

The Pyramid Nile client supports TCP/IP, which comes standard with the system.

Hardware Supported

Any Pyramid Nile or RM1000 system

File System Considerations

ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

SCO

OPEN SYSTEMS SOFTWARE

- **Operating System**
  - ✓ **SCO UNIX 386 3.2.4**
  - ✓ **SCO Open Destop 2.0 or 3.0**
- **Communication Protocols**
  - ✓ **TCP/IP (Standard for Open Desktop)**
  - ✓ **SCO TCP/IP Runtime System 1.1.3**
    (for SCO UNIX)
- **Hardware Supported**
  - ✓ **Any 386SX or Higher PC**
- **File System Considerations**

SCO UNIX 386 Client

The Santa Cruz Operation (SCO) platform is a UNIX-based operating system for Intel processors. It supports a backup-archive client CLI and GUI and an administrative client CLI.

Operating System

The SCO client requires either SCO UNIX 386 3.2.4 or SCO Open Desktop 2.0 or 3.0.

Communication Protocols

The SCO client supports TCP/IP, which comes standard with SCO Open Desktop 2.0 or 3.0. For SCO UNIX 386, TCP/IP is available through the SCO TCP/IP Runtime System 1.1.3.

Hardware Supported

Any 386SX or higher level personal computer or workstation.

File System Considerations

ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

Sequent PTX Client

The Sequent PTX client supports a backup-archive client CLI and GUI and an administrative client CLI.

Operating System

The Sequent PTX client requires Sequent PTX UNIX 2.1.6 or 2.1.7.

Communication Protocols

The Sequent PTX client supports TCP/IP through PTX/TCP/IP 2.4.3, which comes standard with the system.

Hardware Supported

The Sequent PTX client is supported by the following systems:

- Symmetry 5000 System (S5000) Model SE20, SE30, SE60, or SE70

- Symmetry 2000 System (S2000) Model 200, 250, 290, 400, 450, 490, 700, 750, or 790.

File System Considerations

ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

```
┌─────────────────────────────────────────────────────────────────┐
│                        SINIX RISC Client                         │
│  ───────────────────────────────────────────────────────────     │
│                                                                   │
│                          SIEMENS                                  │
│                         ▨▨▨▨▨▨▨▨                                  │
│                          NIXDORF                                  │
│                                                                   │
│              ▪ Operating System                                   │
│                 ✓ Siemens Nixdorf Sinix-N or Sinix-P 5.4.2        │
│              ▪ Communication Protocol                             │
│                 ✓ TCP/IP (Standard)                               │
│              ▪ Hardware Supported                                 │
│                 ✓ Siemens Nixdorf RISC Type RM200,                │
│                    RM400, or RM600                                │
│              ▪ File System Considerations                         │
│                                                                   │
└─────────────────────────────────────────────────────────────────┘
```

Sinix RISC Client

The Siemens Nixdorf Sinix RISC client supports a backup-archive client CLI and GUI and an administrative client CLI.

The Siemens Nixdorf Sinix RISC client also provides the support for Pyramid Nile systems; see "Pyramid Nile Client" on page 304.

Operating System
The Sinix RISC client requires Siemens Nixdorf UNIX Sinix-N or Sinix-P Version 5.4.2.

Communication Protocol
The Sinix RISC client supports TCP/IP, which comes standard with the system.

Hardware Supported
The Sinix RISC client runs on Siemens Nixdorf RISC machines Type RM200, RM400, or RM600.

File System Considerations
ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

SINIX 386/486 Client

SIEMENS

NIXDORF

- ■ **Operating System**
    - ✓ **Siemens Nixdorf Sinix-Z 5.4.21**
- ■ **Communication Protocol**
    - ✓ **TCP/IP (Standard)**
- ■ **Hardware Supported**
    - ✓ **Siemens Nixdorf PCD-4H or Compatible**

      Workstation with 386 or 486 Processor
- ■ **File System Considerations**

Sinix 386/486 Client

The Siemens Nixdorf Sinix 386/486 client supports a backup-archive client CLI and GUI and an administrative client CLI.

Operating System
The Sinix 386/486 client requires the Siemens Nixdorf UNIX Sinix-Z Version 5.4.1.

Communication Protocol
The Sinix 386/486 client supports TCP/IP, which comes standard with the system.

Hardware Supported
The Sinix 386/486 client runs on the Siemens Nixdorf PC Type PCD-4H or compatible workstation with a 386 or 486 processor.

File System Considerations
ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

SGI IRIX Client

```
                    SGI IRIX Client

                    ❋  Silicon Graphics


               ▪ Operating System
                  ✓ IRIX 5.2 or 5.3
               ▪ Communication Protocol
                  ✓ TCP/IP (Standard)
               ▪ Hardware Supported
                  ✓ Silicon Graphics Workstation with
                    MIPS Processor
               ▪ File System Considerations
```

SGI IRIX Client

The Silicon Graphics IRIX UNIX client supports a backup-archive client CLI and GUI and an administrative client CLI. Applications using the ADSM API are also supported.

Operating System
The IRIX client requires the SGI IRIX UNIX Release 5.2, 5.3, or 6.2.

Communication Protocol
The IRIX client supports TCP/IP, which comes standard with the system.

Hardware Supported
The IRIX client runs on the SGI workstations with MIPS processor.

File System Considerations
ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

Sun Solaris/SunOS Client

- **Operating System**
  - ✓ **SunOS 4.1.3**
  - ✓ **Solaris 2.3, 2.4, or 2.5**
- **Communication Protocol**
  - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
  - ✓ **Sun-4 or Compatible Workstation**
  - ✓ **SPARC (Solaris 2.5 only) or Compatible Workstation**
- **File System Considerations**

Sun Solaris Client

The Sun Microsystems SunOS or Solaris clients support backup-archive client and administrative client CLIs and GUIs. Applications using the ADSM API are also supported. The Solaris 2.5 client is provided as a separate PTF.

Space Management is provided for Sun Solaris 2.5.

Operating System
The SunOS client supports SunOS 4.1.3. The Solaris client supports Solaris 2.3, 2.4, or 2.5.

The SunOS and Solaris clients are all supported in native mode, but for the administrative GUI, Solaris 2.3 and 2.4 requires the binary compatibility packages. Binary compatibility means that the application runs under Solaris but uses only the Sun-4 architecture. For more information about the binary compatibility packages, refer to your Solaris documentation. Note that Solaris 2.4 and 2.5 are supported only on the SPARC workstations.

An X Windows system such as OpenWindows 2.0, which comes standard with SunOS or Solaris, is required for the backup-archive GUI. The administrative GUI will work with the OpenLook Window Manager, but the Motif Window Manager is recommended.

The Sun Solaris 2.5 HSM client requires Sun Solaris 2.5 and the Common Desktop Environment, which comes standard with Solaris 2.5.

Communication Protocol

The SunOS or Solaris clients support TCP/IP, which comes standard with SunOS or Solaris.

Hardware Supported

The SunOS or Solaris clients will work on any Sun-4, SPARC, or compatible workstation. The Sun Solaris 2.5 HSM client requires a single-processor Sun SPARC workstation. There is no support for MP.

File System Considerations

ADSM provides backup and archive services for the following file systems:

- All locally mounted file systems
- NFS

ADSM backs up regular files and directories. It does not back up character special files, block special files, or FIFO special files. When you back up your files, ADSM also saves standard UNIX access permissions assigned to the files.

Windows Client

Windows Client

■ **Communication Protocols**
  - ✓ **TCP/IP**
    FTP Software PC/TCP 3.0 for DOS
    IBM TCP/IP 2.1 for DOS
    NetManage Chameleon 4.0 for Windows
    Novell LAN Workplace 4.1 and 4.2 for DOS
    Trumpet TCP/IP 1.0 Rev. A
    WollonGong PathWay Runtime 2.0 for
    DOS or Windows
    Microsoft TCP/IP for Windows 3.11 for
    Workgroups
  - ✓ **IPX/SPX**
    NetWare Workstation for Windows
  - ✓ **IBM NetBIOS**
    IBM LAN Support Program
  - ✓ **APPC**
    Networking Services/Windows 1.0
    IBM Client Access/400 for Windows

■ **Operating System**
  - ✓ **Microsoft Windows V3.1, 3.11**
  - ✓ **Windows for WorkGroups 3.11**

■ **Hardware Supported**
  - ✓ **Any PC That Supports Windows**

Windows Client

The Microsoft Windows client supports backup-archive client and administrative client GUIs. An administrative client CLI is supported through a pseudo-text-mode window application that simulates a command line inside a window. Applications using the ADSM API are also supported.

Operating System

The Windows client requires Microsoft Windows 3.1 or Windows for Workgroup 3.11.

Communication Protocols

The Windows client supports any of the following protocols:

- TCP/IP, which is available through one of these WinSock TCP/IP stacks:

    - FTP Software PC/TCP 3.0 for DOS

    - IBM TCP/IP 2.1 for DOS

    - NetManage Chameleon 4.0 for Windows

    - Novell LAN Workplace 4.1 or 4.2 for DOS

    - Trumpet TCP/IP 1.0 Revision A

    - WollonGong PathWay Runtime 2.0 for DOS or Windows

    - Microsoft TCP/IP for Windows 3.11 for Workgroups

- IPX/SPX, which is available through NetWare Workstation for Windows

- NetBIOS, which is available through the IBM LAN Support Program

- APPC, which is available through Networking Services/Windows 1.0 or IBM Client Access/400 for Windows.

- 3270, which is available through IBM Personal Communication/3270 3.0 or later

Hardware Supported
Any personal computer with at least 4MB of memory

File System Considerations
The FAT is a hierarchical file system of directories, subdirectories, and files. With DOS being a single-user system, the files have no concept of ownership.

Windows NT 3.5 or Later
Windows 95
- **Communication Protocols**
  - ✓ **TCP/IP**
  - ✓ **IPX/SPX**
  - ✓ **NetBIOS**
- **Hardware Supported**
  - ✓ **Any Intel-based 80386 or Greater**
- **File System Considerations**

Windows 32 bit Client

The Microsoft Windows 32 bit client was formerly referred to as the Windows NT client. It supports a backup-archive client GUI, and and administrative client CLI and GUI and runs on Windows NT and Windows 95. . Applications using the ADSM API are also supported.

Operating System

The Windows 32 bit client requires Microsoft Windows NT Version 3.5 or later, or Windows 95.

Communication Protocols

The Windows 32 bit client supports the following protocols:

- TCP/IP

- IPX/SPX

- NetBIOS

all of which come standard with the operating system.

Hardware Supported

Any Intel-based 80386 processor or greater. For Windows 95, at least 8 MB of memory is recommended; for NT/Workstation, a minimum of 12 MB; and for NT/Server, at least 32 MB.

File System Considerations

The ADSM Windows 32 bit client supports the FAT file system, the HPFS from OS/2, and its own special NT File System (NTFS). NTFS is similar to FAT and HPFS and provides a hierarchy of files and directories.

- FAT

    FAT support on NT is the same as on DOS except NT supports long and mixed-case file names.

- HPFS

    Backing up EA data is not supported. Microsoft is dropping support for HPFS in NT 4.0, so there are no plans to add this support in ADSM.

- NTFS

    ADSM provides full support for NTFS file security. All NTFS file level audit and permission information can be backed up with the ADSM Windows 32 bit client running on Windows NT. Every NTFS file and directory has four security descriptors that ADSM backs up: owner security identifier (SID), primary group SID, discretionary ACL, and system ACL.

    Backup of Macintosh files or folders on NTAS with File Services for Macintosh running is not supported.

    The POSIX file system is not supported for ADSM backup.

Windows NT

On Windows NT, special commands are provided to backup the NT registry, on both the CLI and GUI. The registry stores information about hardware, applications, and operating system settings. It contains similar information as the WIN.INI and SYSTEM.INI files in Windows 3.1 and the CONFIG.SYS and AUTOEXEC.BAT files in OS/2. Applications designed for Windows NT also store their own .INI settings in the registry.

The registry is automatically backed up during an ADSM incremental backup of the system drive when the BACKUPREG option is set. This option is set by default.

ADSM backs up the registry in a two-stage process. First, ADSM uses an NT API to back up the registry into a directory structure on the system drive under a root directory name of ADSM.SYS. Then ADSM backs up the directories and files under ADSM.SYS to an ADSM server.

When you use ADSM to restore the registry, ADSM restores the local registry backup in ADSM.SYS if it exists. If there is no backup in ADSM.SYS, ADSM restores the registry from the ADSM server.

Windows 95

There are some limitations when running the Windows 32 bit client on a Windows 95 system. The backup-restore registry function is currently not implemented for Windows 95. Also, backing up remote NTFS drives from Windows 95 is not recommended, as none of the NTFS security information will be backed up. This is the case because the Windows 32 bit client APIs dealing with security are available only on Windows NT.

AS/400 API Support
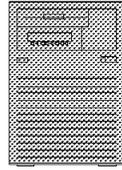
- **Operating System**
  - ✓ **OS/400 V3R1 or V3R2 (CISC Systems)**
  - ✓ **OS/400 V3R6 or V3R7 (RISC Systems)**
- **Communication Protocols**
  - ✓ **TCP/IP (Standard)**
  - ✓ **APPC (Standard)**
- **Hardware Supported**
  - ✓ **Any AS/400 Supporting the Required OS/400**

AS/400 API support

ADSM for AS/400 provides an AS/400 API that enables users to develop their own application so that their AS/400 system functions as an ADSM backup-archive client.

Operating System
The AS/400 API requires either OS/400 V3R1 or V3R2 for CISC systems or OS/400 V3R6 or V3R7 for RISC systems.

An Integrated Language Environment (ILE) is required to use the AS/400 API. The first ILE compiler available is ILE C/400.

Communication Protocols
The AS/400 API supports either TCP/IP or APPC (LU6.2), both of which come standard with OS/400.

Hardware Supported
Any AS/400 system supports the AS/400 API

CRAY UNICOS Client

- **Operating System**
  - ✓ **CRAY UNICOS 8 or 9**
- **Communication Protocol**
  - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
  - ✓ **CRAY SMP models J90, C90, YMP, and T90**

CRAY UNICOS Client

The CRAY UNICOS client is a service offering available from Storage Manufacturing Information Technology (SM IT) in Mainz, Germany. A backup-archive client CLI is supported. Refer to "ADSM CRAY UNICOS Client" on page 172 for further details.

Operating System
The CRAY UNICOS client requires UNICOS 8 or 9.

Communication Protocol
The CRAY UNICOS client supports TCP/IP, which comes standard with the operating system.

Hardware Supported
The CRAY SMP models J90, C90, YMP, and T90 are supported.

Digital OpenVMS Client

d i g i t a l
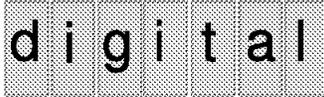
- **Operating System**
  - ✓ **OpenVMS V5.5-2 or Higher (VAX Systems)**
  - ✓ **OpenVMS V6.1 or Higher (Alpha Systems)**
- **Communication Protocol**
  - ✓ **TCP/IP (through DEC TCP/IP Services for OpenVMS)**
- **Hardware Supported**
  - ✓ **Any VAX or Alpha System Supporting the Required Operating System**

Digital OpenVMS Client

The Digital OpenVMS client is an offering available from Storage Solution Specialist Inc. A backup-archive client CLI is supported. The client support is for ADSM API V2.1.2 or higher. Refer to "OpenVMS Client for ADSM" on page 176 for further details.

Operating System
The OpenVMS client requires OpenVMS V5.5-2 or higher for VAX systems, and OpenVMS V6.1 or higher for Alpha systems.

Communication Protocols
The OpenVMS client supports TCP/IP, which is available through the DEC TCP/IP Services for OpenVMS.

Hardware Supported
Any valid VAX or Alpha system supporting the operating system environment

Fujitsu UXP Client
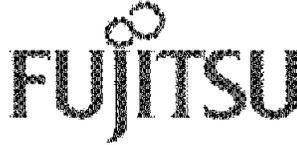
Operating System

✓ **UXP/P for VPP Series**

✓ **UVP/M for VPS600 Series**

Communication Protocol

✓ **TCP/IP (Standard)**

Hardware Supported

✓ **Any VPP300 or VPS600 Model**

Fujitsu UXP/V Client

The Fujitsu UXP/V client is a service offering available from SM IT in Mainz, Germany. A backup-archive and administrative client CLI is supported. Applications using the ADSM API are also supported. Refer to "ADSM Fujitsu UXP/V Client" on page 174 for further details.

Operating System

The Fujitsu UXP/V client requires the Fujitsu UXP/P operating system for Fujitsu VPP300 series or the UVP/M for the VPS600 series.

Communication Protocol

The Fujitsu UPX/V supports TCP/IP, which is available through the operating system.

Hardware Supported

Any models of the Fujitsu VPP300 series or VPS600 series

- **Operating System**
  - ✓ **OS/2 Warp 3.0 or Warp Server**
- **Communication Protocols**
  - ✓ **Refer to the OS/2 Client**
- **Hardware Supported**
  - ✓ **Refer to the OS/2 Client**

Lotus Notes Backup Agent

The ADSM Lotus Notes Backup Agent is an integrated Lotus Notes application running on OS/2, with CLI and GUI function. It functions as a Lotus Notes client on the OS/2 supported machines and performs incremental and full backups of OS/2 Notes databases at the document level. The Agent uses the Lotus Notes and ADSM API to back up and restore OS/2 notes documents and databases to and from an ADSM server.

For further details see "ADSM and Lotus Notes" on page 240, and refer to the redbook entitled *Using ADSM to Back Up Lotus Notes* (SG24-4534).

Operating System
The Lotus Notes Backup Agent requires the ADSM OS/2 API, which is available through OS/2 Warp 3.0 or Warp Server.

Communication Protocols
The ADSM Lotus Notes Backup Agent uses the communication support of OS/2.

Hardware Supported
Any OS/2 Warp 3.0 or Warp Server supported hardware

Tandem Guardian Client

TANDEM

- **Operating System**
  - ✓ **Tandem Guardian NSK Version D30 and Higher**
- **Communication Protocol**
  - ✓ **TCP/IP (Standard)**
- **Hardware Supported**
  - ✓ **Any Tandem Model that Supports the Required Operating System**

Tandem Guardian Client

The Tandem Guardian platform is supported by the Backhome! client provided by ETI, Montreal, Canada. A backup-archive client CLI is supported. Refer to "Tandem Guardian Client" on page 178 for further details.

Operating System

The Backhome! client requires Tandem Guardian (NSK) Version D30 or higher.

Communication Protocol

The Backhome! client supports TCP/IP, which is available through the operating system.

Hardware Supported

Any Tandem model that supports the Version D30 operating system

---

WebShell Interface

```
http://WWW.ADSM.CLIENT
```

- **Operating System**
  - ✓ **Refer to the Required Operating Systems for ADSM**
    - **OS/2, Windows 95, Windows NT, and AIX clients**
- **Communication Protocol**
  - ✓ **Internet Access**
- **Hardware Supported**
  - ✓ **Any Machine That Supports a Web Browser**

WebShell Interface

The WebShell interface can be used to manage backup and restore operations from any machine, on any platform, on any network. With it, an ADSM client can back up and restore its own data, or an ADSM administrator can centralize the backup and restore operations of many clients. Clients with the WebShell installed can be accessed from Web browsers running on any other system. For more information refer to Chapter 5, "WebShell Client" on page 111.

Operating System

The WebShell interface is packaged with the following clients: ADSM OS/2, Windows 95, Windows NT, and AIX.

A Web browser that supports HTML 2.0 or higher is required for accessing the WedShell interface. Examples of such browsers are Web Explorer, Netscape, and Microsoft Internet Explorer.

Communication Protocol

Access to the Internet is required.

Hardware Supported

At least one ADSM server must be available.

Appendix C. Device Support

ADSM Device Support

Basic License
– Disk Devices
– Single Unit Tape and Optical Devices

Device Support Module 1
– Higher Capacity Tape Devices
– Low Capacity Tape and Optical Libraries

Device Support Module 2
– 1/2 Inch Tape Devices
– Medium Capacity Tape and Optical Libraries

Device Support Module 3
– Medium Capacity 1/2 Inch Tape Libraries
– High Capacity Tape and Optical Libraries

Device Support Module 4
– Very High Capacity 1/2 Inch Tape Libraries
– Very High Capacity DLT and Optical Libraries

ADSM Device Support

With any ADSM server license there is a base device support that will support all basic disk and tape devices in a system. Further device support can be purchased through Device Support Modules 1 through 4. With this modularity, you only pay for the device Support needed. The Device Support Modules cover different devices for the different ADSM servers, but in general terms the device support can be grouped as indicated. Some of the tape and optical devices are available only for certain server platforms.

A higher level Device Support Module always includes support for the devices on the lower levels, that is, the modules are cumulativ e. Most hardware companies do not provide their own device drivers, which is why ADSM provides device drivers. For the IBM 349X tape drives and libraries, as well as the IBM 3590, ADSM uses the drivers provided with the hardware.

The devices supported for the different platforms are constantly being updated. Device support flyers are referenced in the description of each platform, and these are available through IBM. For the most current list, look up the ADSM home page on the Web, or the ADSM family page and the corresponding device support pages for each server:

- http://www.storage.ibm.com/adsm

- http://www.storage.ibm.com/storage/software/adsm/ads0fam.htm

Base License

The basic ADSM license supports all disk devices in the system, as well as the most general single tape devices for the specific platform. The following device classes are typically included in the base support:

• All disk devices supported by the hardware and operating system

• Single unit 4 mm tapes (2 - 5 GB)

• Single unit 8 mm tapes (2.3 -10 GB)

• QIC tapes (1.2 GB)

• Single unit optical drives (up to 14 GB)

Device Support Module 1

Device Support Module 1 includes all devices supported by the base license. In addition it supports higher capacity single units, and the most common low capacity tape libraries (4 and 8 mm autochangers) or optical libraries (jukeboxes). The following device classes are grouped in Support Module 1:

• Single unit DLT (10 - 20 GB)

• 4 mm tape libraries (up to 68 GB)

• 8 m tape libraries (up to 70 GB)

• Optical libraries (16 slots)

Device Support Module 2

Device Support Module 2 includes all devices supported by Module 1. In addition it supports 1/2 inch tape transports and medium capacity tape libraries (DLT, 4 and 8 mm autochangers) or optical libraries (jukeboxes). The following device classes are grouped in Support Module 2:

• Single unit 1/2 inch tapes (0.2 - 10 GB tape cassettes)

• 4 mm tape libraries (up to 160 GB)

• 8 mm tape libraries (up to 380 GB)

• DLT libraries (up to 280 GB)

• Optical libraries (up to 52 slots)

The ADSM servers for MVS, VM, AS/400, and VSE/ESA include Support Modules 1 and 2 with the base license. For these servers the 4 and 8 mm tapes, DLT, and optical devices are generally not supported.

Device Support Module 3

Device Support Module 3 includes all devices supported by Module 2. In addition it supports 1/2 inch tape libraries and high capacity tape libraries (DLT, 4 and 8 mm autochangers) or optical libraries (jukeboxes). The following device classes are grouped in Support Module 3:

• 1/2 inch tape 3590 in random mode (10 - 100 GB)

• 1/2 inch tape IBM 3494 (240 slots)

• 4 mm tape libraries (up to 240 GB)

• 8 mm tape libraries (up to 580 GB)

- DLT libraries (up to 600 GB)

- Optical libraries (up to 170 slots)

The ADSM server for VSE/ESA currently has no support for devices in Module 3.

Device Support Module 4

Device Support Module 4 includes all devices supported by Module 3. In addition it supports very high capacity 1/2 inch tape libraries and very high capacity DLT and optical libraries. The following device classes are grouped in Support Module 4:

- 1/2 inch tape IBM 3494 (up to 3,040 slots)

- 1/2 inch tape IBM 3495 (up to 100,000 slots)

- Other 1/2 inch tape libraries

- DLT libraries (up to 2.5 TB)

- Optical libraries (up to 1054 slots)

The ADSM servers for VSE/ESA, OS/2, and HP-UX currently do not support devices in Module 4.

ADSM Device Support Units

Disks		4 mm Tape Autochangers
4 mm Tape Drives		8 mm Tape Autochangers
8 mm Tape Drives		
QIC Tape Drives		DLT Libraries
DLT Drives		
DTF Drives		1/2" Tape Libraries
1/2" Tape Drives		
Optical Drives		Optical Libraries

ADSM Device Support: Units

Any ADSM server supports all standard disk devices available for that server. Such support is provided through the operating system's device support routines. The different servers have a wide variety of supported storage devices. This is dependent on the device attachments provided by the server, and the most common method for device attachment is the SCSI interface. For personal computers, ISA/EISA and PSI are also common attachments, and for mainframe systems, different types of channel attachments are provided.

Device Types

The graphic shows the different device types that are supported by one more servers. They include 4 and 8 mm single tapes and tape libraries (or autochangers), 1/4″ QIC tapes and 1/2″ tape reels, DTF tapes and DLT tapes and libraries, optical devices and optical device jukeboxes, and 1/2″ cassette tapes and libraries. The device support is continuously updated as new devices are added. The list of supported devices will be different for each server. For the most up-to-date information about devices supported with each server, consult the Web pages.

There is also a set of device support flyers, one for each server. These are continuously updated and can be ordered through IBM. Call the IBM Information Support Center in the United States at 1-800-IBM-2468 and ask for "STAR 20," or contact your local IBM representative or reseller. The flyers can also be ordered through the IBM FAX Information Service, fax number 1-800-IBM-4FAX inside the United States and Canada, or fax number 1-415-855-4329 outside the United States. The flyer entitled *ADSM for AIX Supported Storage Devices* (G225-6620) is continuously updated and shows the latest supported devices.

- ADSM for AIX Supported Storage Devices - G225-6620
- ADSM for HP-UX Device Validation List - G225-6627
- ADSM for MVS Supported Storage Devices - G225-6619
- ADSM for Sun Solaris Device Validation List - G225-6626
- ADSM for VM Device Support List - G225-6624

- ADSM for VSE/ESA Device Validation List - G225-6623
- ADSM for Windows NT Supported Storage Devices - G225-6672
- ADSM/2 Device Support List - G225-6621
- ADSM/400 Supported Storage Devices - G225-6619
- ADSM/400 2.3 Device Validation List - G225-6625

Appendix D. Special Notices

This publication is intended to help customers, consultants, Business Partners and IBMers to understand ADSM concepts and to learn what is new with ADSM. The information in this publication is not intended as the specification of any programming interfaces that are provided by ADSM, Lotus Notes, SQL-BackTrack, AvailHSM for NetWare or OPEN/Stor for Windows NT. See the PUBLICATIONS section of the IBM Programming Announcement for ADSM for AIX, ADSM for HP-UX, ADSM for Sun Solaris, ADSM for Windows NT, ADSM for MVS, ADSM for VM, ADSM for OS/400, ADSM for VSE/ESA, and ADSM for OS/2 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (″vendor″) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these

names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ACF/VTAM	AD/Cycle
ADSTAR	AIX
AIX/6000	Application System/400
AS/400	AT
BookManager	Business Partner
C/370	C/400
Client Access/400	DB2
DB2/2	DB2/6000
DFSMS	DFSMS/MVS
DFSMShsm	Enterprise Systems Architecture/390
ESCON	HACMP/6000
IBM Business Partner	IBM
IBMLink	Integrated Language Environment
Language Environment	Magstar
MVS/DFP	MVS/ESA
NetView	OpenEdition
Operating System/2	Operating System/400
OS/2	OS/390
OS/400	PowerPC
PROFS	RACF
RAMAC	RS/6000
S/390	SAA
Scalable POWERparallel Systems	SP
SP2	StorageSmart
System/390	SystemView
VM/ESA	VSE/ESA
VTAM	Workplace

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Andataco	Andataco Corporation
Apple	Apple Computer, Inc.
Attachmate	Attachmate Corporation
AvailHSM, OPEN/stor	Wang Software Storage Management
CompuServe	CompuServe, Inc.
DECstation	Digital Equipment Corporation
DLT	Quantum Corporation

DPX/20	Groupe Bull
Exabyte	Exabyte Corporation
Extra!	Attachmate Corporation
Hewlett-Packard, HP-UX	Hewlett-Packard Company
Intel	Intel Corporation
IPX/SPX, NetWare, Novell	Novell, Inc.
IRIX	Silicon Graphics, Inc.
1-2-3, Lotus, Freelance, Freelance Graphics	Lotus Development Corporation
Macintosh	Apple Computer, Inc.
Motif	Open Software Foundation, Inc.
Open Desktop	The Santa Cruz Operation, Inc.
Sequent	Sequent Computer Systems
Sinix	Siemens Nixdorf Information Systems, Inc.
Sun, SunOS, SPARCstation, Network File System, NFS	Sun Microsystems, Inc.
StorageTek	Storage Technology Corporation
ULTRIX	Digital Equipment Corporation
WangDAT	Tecmar Technologies, Inc.
Windows 95, Windows NT	Microsoft Corporation
X Windows	Massachusetts Institute of Technology

Other trademarks are trademarks of their respective companies.

Appendix E. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 337.

ADSM Redbooks

General Topics

Book Title	Publication Number
General Topics	
ADSM Concepts	SG24-4877
ADSM Version 2 Presentation Guide	SG24-4532
ADSM Advanced Implementation Experiences	GG24-4221
Using ADSM Hierarchical Storage Management	SG24-4631
Client Disaster Recovery: Bare Metal Restore	SG24-4880
Specific Server Books	
ADSM Server for Windows NT Configuration and Recovery Examples	SG24-4878
Getting Started with ADSM/6000	GG24-4421
ADSM for AIX: Advanced Topics	SG24-4601
AIX Tape Management	SG24-4705
ADSM/6000 on 9076 SP2	GG24-4499
ADSM for MVS: Recovery and Disaster Recovery	SG24-4537
ADSM for MVS: Using Tapes and Tape Libraries	SG24-4538
Getting Started with ADSM/2	GG24-4321
ADSM for OS/2: Advanced Topics	SG24-4740
Setting Up and Implementing ADSM/400	GG24-4460
ADSM/VSE Implementation Guide	SG24-4266
Specific Client Books	
Getting Started with ADSM NetWare Clients	GG24-4242
Getting Started with ADSM AIX Clients	GG24-4243
ADSM API Examples for OS/2 and Windows	SG24-2588
ADSM with Other Products	
Using ADSM to Back Up Databases	SG24-4335
Using ADSM to Back Up Lotus Notes	SG24-4534
HSM for NetWare: ADSM and AvailHSM Implementation	SG24-4713
Using ADSM to Back Up OS/2 LAN Server & Warp Server	SG24-4682
Backup, Recovery, and Availability with DB2 PE	SG24-4695

Other Redbooks

Book Title	Publication Number
TME 10 Cookbook for AIX: Systems Management and Networking	SG24-4867
Planning for IBM Remote Copy	SG24-2595
P/DAS and Enhancement to the IBM 3990-6 and RAMAC Array Family	SG24-4724
Backup, Recovery, and Availability with DB2 Parallel Edition on RS/6000 SP	SG24-4695

Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

ADSM Product Publications

Book Title	Publication Number
ADSM General Information	GH35-0131
ADSM V2 Installing the AIX Server and Administrative Client	SH35-0136
ADSM V2 for AS/400 Quick Start	GA32-0357
ADSM V2 for HP-UX Quick Start	GC35-0256
ADSM V2 Installing the MVS Server and Administrative Client	SH26-4043
ADSM V2 for OS/2 Quick Start	GC35-0231
ADSM V2 for Sun Solaris Quick Start	GC35-0262
ADSM V2 for VM Quick Start	GC35-0227
ADSM V2 for Windows NT Quick Start	GC35-0235
ADSM V2 for AIX Administrator's Guide	SH35-0134
ADSM V2 for AS/400 Administrator's Guide	SC35-0196
ADSM V2 for HP-UX Administrator's Guide	GC35-0257
ADSM V2 for OS/2 Administrator's Guide	GC35-0232
ADSM V2 for Sun Solaris Administrator's Guide	GC35-0263
ADSM V2 for MVS Administrator's Guide	SH26-4039
ADSM V2 for VM Administrator's Guide	GC35-0228
ADSM V2 for Windows NT Administrator's Guide	GC35-0236
ADSM V2 for AIX Administrator's Reference	SH35-0135
ADSM V2 for AS/400 Administrator's Reference	SC35-0197
ADSM V2 for HP-UX Administrator's Reference	GC35-0258
ADSM V2 for MVS Administrator's Reference	SH26-4040
ADSM V2 for OS/2 Administrator's Reference	GC34-0233
ADSM V2 for Sun Solaris Administrator's Reference	GC35-0264
ADSM V2 for VM Administrator's Reference	GC35-0229
ADSM V2 for Windows NT Administrator's Reference	GC35-0237
ADSM V2 for MVS DRM Administrator's Guide and Reference	GC35-0238
ADSM V2 Messages	SH35-0133
ADSM V2 Device Configuration	SH35-0137
ADSM V2 Installing the Clients	SH26-4049
ADSM V2 AFS/DFS Backup Clients	SH26-4048
ADSM V2 Using the UNIX HSM Clients	SH26-4030
ADSM V2 Using the UNIX Backup-Archive Client	SH26-4052
ADSM V2 Using the OS/2 Backup-Archive Client	SH26-4053
ADSM V2 Using the DOS Backup-Archive Client	SH26-4054
ADSM V2 Using the Microsoft Windows Backup-Archive Clients	SH26-4056
ADSM V2 Using the Novell NetWare Backup-Archive Client	SH26-4055
ADSM V2 Using the Apple Macintosh Backup-Archive Client	SH26-4051
ADSM V2 Using the Lotus Notes Backup Agent	SH26-4047
ADSM V2 Using the Application Programming Interface	SH26-4002

Book Title	Publication Number
ADSM V2 Reference Cards for the Backup-Archive Clients	SX26-6013

ADSM Online Product Library

All of the ADSM publications are available in online readable format on the CD-ROM listed below. The ADSM library is also available on the following CD-ROMs: These books can also be ordered in softcopy format on CD-ROM:

CD-ROM Title	Publication Number
ADSM Online Product Library	SK2T-1893
MVS Base Collection Kit	SK2T-0710
VM Base Collection Kit	SK2T-2067
AS/400 Base Collection Kit	SK2T-2171

Other Publications

Book Title	Publication Number
Backup/Restore Interface for ORACLE for SAP R/3	SC33-6379
TME 10 Tivoli/Plus for ADSM User's Guide	GC31-8405
The DFS for AIX/6000	GG24-4255
ADSM AFS/DFS Backup Clients V2.1	SH26-4048
IBM 3466 Network Storage Manager Introduction and Planning Guide	GA32-0374
IBM 3466 Network Storage Manager Start Up Guide	GA32-0373
CLIO/S V2.2 User's Guide and Reference	GC28-1676
OnDemand for AIX: Introduction and Planning Guide	G544-5281
Remote Copy Administrator's Guide and Reference	SC35-0169

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL http://www.redbooks.ibm.com.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

      TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
      TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)

  To get lists of redbooks:

      TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
      TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE

  To register for information on workshops, residencies, and redbooks:

      TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996

  For a list of product area specialists in the ITSO:

      TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE

- **Redbooks Home Page on the World Wide Web**

  http://w3.itso.ibm.com/redbooks

- **IBM Direct Publications Catalog on the World Wide Web**

  http://www.elink.ibmlink.ibm.com/pbl/pbl

  IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU

Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
(+45) 48 14 2207 (long distance charge)	Outside North America

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com

- **On the World Wide Web**

Redbooks Home Page	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank).

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name _____ Last name _____

Company _____

Address _____

City _____ Postal code _____ Country _____

Telephone number _____ Telefax number _____ VAT number _____

- Invoice to customer number _____

- Credit card number _____

Credit card expiration date _____ Card issued to _____ Signature _____

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

Index

Numerics

2000 (year)   140, 246
2GB files   144
3466
   benefits   127
   description   125
   further customization   129
   initial configuration   129
   installing   129
   installing the clients   131
   overview   69
   package   125
   WebShell client   131
3494   125, 194, 324
3495   194
3570   168
3590   125, 324
3990-6   192
4 mm tape libraries   324
400COMMBUFFERBYTES   151
400SENDBUFFERBYTES   151
4GB file spaces   144
8 mm tape libraries   324

A

Adabas   60
ADINT/ADSM   60, 238, 239
administrative client   41
administrative GUI   42, 44
administrator
   authority   42
   CLI   42
   features   42
   functions   41
   GUI   42, 44
   span of control   43
   support   42
   WebShell   131
ADSM
   complementary products   171
   CompuServe forum   65
   discussion list   65
   FTP server   65
   homepage   65
   implementation   68
   information sources   65
   paradigm   54
   publications   333
   redbooks   65, 67, 333
   services   68
   updates   65
   web access   65

ADSM administrator
   definition   12
   description   14
ADSM client
   definition   12
ADSM components   12
ADSM consultant   65, 66
ADSM database
   alternate volumes   39
   backup   35, 39, 75
   description   38
   extension   39
   full backup   39
   incremental backup   39
   mirroring   39
   recovery log   39
   restoring from disaster   81
ADSM homepage   65
ADSM macros   80
ADSM server
   definition   12
   description   23
   supported platforms   11
ADSM-L discussion list   66
ADSMConnect Agent for Oracle Backup   155
ADSMPIPE   208
adstarsw@vnet.ibm.com   65
AFS and DFS
   backup   107, 108
   buta   108
   butc   107
   cache manager   104, 105
   comparison   106
   description   103
   file backup support   109
   include/exclude   109
   mount point backup   109
   POSIX compliance   105
   support   103
   token manager   105
   virtual mount points   110
   volume dump support   109
AIX bootable image   136
AIX client   285
analyst authority   42
API
   ADSMConnect Agent use of   221
   applications using   58, 62
   AS/400   153, 316
   AvailHSM use of   182
   BACKINT/ADSM use of   239
   buta use of   109
   DB2/6000 use of   212
   description   10

D

K

Kerberos   30, 104

L

labeling   120
LAN Services   190
LANRES   190
Lotus Notes   61, 155, 198, 240—243
Lotus Notes client   320

M

machine recovery   79
Macintosh PowerPC   153
managed nodes   92
management class   24, 25
manual inventory   8
MESSAGEFORMAT   159
MIGFILEEXPIRATION   152
migration
   management class controls   27
   multiple migration processes   159
   selective   27
migration eligibility   27
mirroring   39, 74
mksysb   136
mount point backup   109
MOUNTABLE   82

N

NCR client   296
NEC client   297
NETBIOSBUFFERSIZE   151
NetTAPE   123
   functions   119
   overview   117
NetWare client   298
NetWare HSM   182
network
   capacity   56
   tuning   56
Network Storage Manager
   benefits   127
   description   125
   further customization   129
   initial configuration   129
   installing   129
   installing the clients   131
   overview   69
   package   125
   WebShell client   131
NOAUDITSTORAGE   165
nodename   143
Notes agent   61
NOTMOUNTABLE   82

O

OBackup   155, 221
offline backup   58
offsite backup   35
offsite media management   78
offsite storage pool   35
ON-Bar   198
OnDemand   62, 186
online backup   58
ONSITERETRIEVE   83
Open Systems Environment Support   110
OPEN/stor   19, 184
OpenEdition MVS client   154, 300
OpenVMS   62, 176
operating system level
   AIX client   285
   AIX server   247
   Apple Macintosh client   287
   AS/400 API client   316
   AS/400 server   251
   Auspex client   288
   Bull DPX/2 client   289
   CMS client   290
   CRAY UNICOS client   317
   Digital OpenVMS client   318
   Digital Ultrix client   291
   Digital UNIX client   292
   DOS UNIX client   293
   Fujitsu UXP/V client   319
   HP-UX client   295
   HP-UX server   255
   Lotus Notes client   320
   MVS server   259
   NCR UNIX SVR4 UNIX client   296
   NEC EWS-UX/V client   297
   NetWare client   298
   OpenEdition MVS client   300
   OS/2 client   302
   OS/2 server   263
   Pyramid Nile client   304
   SCO UNIX client   305
   Sequent PTX client   306
   SGI IRIX client   309
   Sinix 386/486 client   308
   Sinix RISC client   307
   Sun Solaris client   310
   Sun Solaris server   267
   Tandem Guardian client   321
   VM server   271
   VSE/ESA server   275
   Windows 95   314
   Windows client   312
   Windows NT client   314
   Windows NT server   279
optical libraries   324
options
   400COMMBUFFERBYTES   151
   400SENDBUFFERBYTES   151

IBM®

Printed in U.S.A.